

Testimony of Ray Kurzweil on the Societal Implications of Nanotechnology

***Committee on Science, U.S. House of Representatives, Hearing
April 9, 2003***

Summary of Testimony:

The size of technology is itself inexorably shrinking. According to my models, both electronic and mechanical technologies are shrinking at a rate of 5.6 per linear dimension per decade. At this rate, most of technology will be “nanotechnology” by the 2020s.

We are immeasurably better off as a result of technology, but there is still a lot of suffering in the world to overcome. We have a moral imperative, therefore, to continue the pursuit of knowledge and advanced technologies, such as nanotechnology, that can continue to overcome human affliction. There is also an economic imperative to continue due to the pervasive acceleration of technology, including miniaturization, in the competitive economy.

Nanotechnology is not a separate field of study that we can simply relinquish. We will have no choice but to confront the challenge of guiding nanotechnology in a constructive direction. There are strategies we can deploy, but there will need to be continual development of defensive strategies.

We can take some level of comfort from our relative success in dealing with one new form of fully non-biological, self-replicating pathogen: the software virus.

The most immediate danger is not self-replicating nanotechnology, but rather self-replicating biotechnology. We need to place a much higher priority on developing vitally needed defensive technologies such as antiviral medications. Keep in mind that a bioterrorist does not need to put his “innovations” through the FDA.

Any broad attempt to relinquish nanotechnology will only push it underground, which would interfere with the benefits while actually making the dangers worse.

Existing regulations on the safety of foods, drugs, and other materials in the environment are sufficient to deal with the near-term applications of nanotechnology, such as nanoparticles.

Full Verbal Testimony:

Chairman Boehlert, distinguished members of the U.S. House of Representatives Committee on Science, and other distinguished guests, I appreciate this opportunity to respond to your questions and concerns on the vital issue of the societal implications of nanotechnology. Our rapidly growing ability to manipulate matter and energy at ever smaller scales promises to transform virtually every sector of society, including health and medicine, manufacturing, electronics and computers, energy, travel, and defense. There will be increasing overlap between nanotechnology and other technologies of increasing influence, such as biotechnology and artificial intelligence. As with any other technological transformation, we will be faced with deeply intertwined promise and peril.

In my brief verbal remarks, I only have time to summarize my conclusions on this complex subject, and I am providing the Committee with an expanded written response that attempts to explain the reasoning behind my views.

Eric Drexler's 1986 thesis developed the concept of building molecule-scale devices using molecular assemblers that would precisely guide chemical reactions. Without going through the history of the controversy surrounding feasibility, it is fair to say that the consensus today is that nano-assembly is indeed feasible, although the most dramatic capabilities are still a couple of decades away.

The concept of nanotechnology today has been expanded to include essentially any technology where the key features are measured in a modest number of nanometers (under 100 by some definitions). By this standard, contemporary electronics has already passed this threshold.

For the past two decades, I have studied technology trends, along with a team of researchers who have assisted me in gathering critical measures of technology in different areas, and I have been developing mathematical models of how technology evolves. Several conclusions from this study have a direct bearing on the issues before this hearing. Technologies, particularly those related to information, develop at an exponential pace, generally doubling in capability and price-performance every year. This observation includes the power of computation, communication – both wired and wireless, DNA sequencing, brain scanning, brain reverse engineering, and the size and scope of human knowledge in general. Of particular relevance to this hearing, the size of technology is itself inexorably shrinking. According to my models, both electronic and mechanical technologies are shrinking at a rate of 5.6 per linear dimension per decade. At this rate, most of technology will be “nanotechnology” by the 2020s.

The golden age of nanotechnology is, therefore, a couple of decades away. This era will bring us the ability to essentially convert software, i.e., information, directly into physical products. We will be able to produce virtually any product for pennies per pound. Computers will have greater computational capacity than the human brain, and we will be completing the reverse engineering of the human brain to reveal the software design of human intelligence. We are already placing devices with narrow intelligence in our bodies for diagnostic and therapeutic purposes. With the advent of nanotechnology, we will be able to keep our bodies and brains in a healthy, optimal state indefinitely. We will have technologies to reverse environmental pollution. Nanotechnology and related advanced technologies of the 2020s will bring us the opportunity to overcome age-old problems, including pollution, poverty, disease, and aging.

We hear increasingly strident voices that object to the intermingling of the so-called natural world with the products of our technology. The increasing intimacy of our human lives with our technology is not a new story, and I would remind the committee that had it not been for the technological advances of the past two centuries, most of us here today would not be here today. Human life expectancy was 37 years in 1800. Most humans at that time lived lives dominated by poverty, intense labor, disease, and misfortune. We are immeasurably better off as a result of technology, but there is still a lot of suffering in the world to overcome. We

have a moral imperative, therefore, to continue the pursuit of knowledge and of advanced technologies that can continue to overcome human affliction.

There is also an economic imperative to continue. Nanotechnology is not a single field of study that we can simply relinquish, as suggested by Bill Joy's essay, "Why the Future Doesn't Need Us." Nanotechnology is advancing on hundreds of fronts, and is an extremely diverse activity. We cannot relinquish its pursuit without essentially relinquishing all of technology, which would require a Brave New World totalitarian scenario, which is inconsistent with the values of our society.

Technology has always been a double-edged sword, and that is certainly true of nanotechnology. The same technology that promises to advance human health and wealth also has the potential for destructive applications. We can see that duality today in biotechnology. The same techniques that could save millions of lives from cancer and disease may also empower a bioterrorist to create a bioengineered pathogen.

A lot of attention has been paid to the problem of self-replicating nanotechnology entities that could essentially form a nonbiological cancer that would threaten the planet. I discuss in my written testimony steps we can take now and in the future to ameliorate these dangers. However, the primary point I would like to make is that we will have no choice but to confront the challenge of guiding nanotechnology in a constructive direction. Any broad attempt to relinquish nanotechnology will only push it underground, which would interfere with the benefits while actually making the dangers worse.

As a test case, we can take a small measure of comfort from how we have dealt with one recent technological challenge. There exists today a new form of fully nonbiological self-replicating entity that didn't exist just a few decades ago: the computer virus. When this form of destructive intruder first appeared, strong concerns were voiced that as they became more sophisticated, software pathogens had the potential to destroy the computer network medium they live in. Yet the "immune system" that has evolved in response to this challenge has been largely effective. Although destructive self-replicating software entities do cause damage from time to time, the injury is but a small fraction of the benefit we receive from the computers and communication links that harbor them. No one would suggest we do

away with computers, local area networks, and the Internet because of software viruses.

One might counter that computer viruses do not have the lethal potential of biological viruses or of destructive nanotechnology. This is not always the case: we rely on software to monitor patients in critical care units, to fly and land airplanes, to guide intelligent weapons in our current campaign in Iraq, and other “mission critical” tasks. To the extent that this is true, however, this observation only strengthens my argument. The fact that computer viruses are not usually deadly to humans only means that more people are willing to create and release them. It also means that our response to the danger is that much less intense. Conversely, when it comes to self-replicating entities that are potentially lethal on a large scale, our response on all levels will be vastly more serious, as we have seen since 9-11.

I would describe our response to software pathogens as effective and successful. Although they remain (and always will remain) a concern, the danger remains at a nuisance level. Keep in mind that this success is in an industry in which there is no regulation, and no certification for practitioners. This largely unregulated industry is also enormously productive. One could argue that it has contributed more to our technological and economic progress than any other enterprise in human history.

Some of the concerns that have been raised, such as Bill Joy’s article, are effective because they paint a picture of future dangers as if they were released on today’s unprepared world. The reality is that the sophistication and power of our defensive technologies and knowledge will grow along with the dangers.

The challenge most immediately in front of us is not self-replicating nanotechnology, but rather self-replicating biotechnology. The next two decades will be the golden age of biotechnology, whereas the comparable era for nanotechnology will follow in the 2020s and beyond. We are now in the early stages of a transforming technology based on the intersection of biology and information science. We are learning the “software” methods of life and disease processes. By reprogramming the information processes that lead to and encourage disease and aging, we will have the ability to

overcome these afflictions. However, the same knowledge can also empower a terrorist to create a bioengineered pathogen.

As we compare the success we have had in controlling engineered software viruses to the coming challenge of controlling engineered biological viruses, we are struck with one salient difference. As I noted, the software industry is almost completely unregulated. The same is obviously not the case for biotechnology. A bioterrorist does not need to put his “innovations” through the FDA. However, we do require the scientists developing the defensive technologies to follow the existing regulations, which slow down the innovation process at every step. Moreover, it is impossible, under existing regulations and ethical standards, to test defenses to bioterrorist agents on humans. There is already extensive discussion to modify these regulations to allow for animal models and simulations to replace infeasible human trials. This will be necessary, but I believe we will need to go beyond these steps to accelerate the development of vitally needed defensive technologies.

With the human genome project, 3 to 5 percent of the budgets were devoted to the ethical, legal, and social implications (ELSI) of the technology. A similar commitment for nanotechnology would be appropriate and constructive.

Near-term applications of nanotechnology are far more limited in their benefits as well as more benign in their potential dangers. These include developments in the materials area involving the addition of particles with multi-nanometer features to plastics, textiles, and other products. These have perhaps the greatest potential in the area of pharmaceutical development by allowing new strategies for highly targeted drugs that perform their intended function and reach the appropriate tissues, while minimizing side effects. This development is not qualitatively different than what we have been doing for decades in that many new materials involve constituent particles that are novel and of a similar physical scale. The emerging nanoparticle technology provides more precise control, but the idea of introducing new nonbiological materials into the environment is hardly a new phenomenon. We cannot say a priori that all nanoengineered particles are safe, nor would it be appropriate to deem them necessarily unsafe. Environmental tests thus far have not shown reasons for undue concern, and it is my view that existing regulations on the safety of foods,

drugs, and other materials in the environment are sufficient to deal with these near-term applications.

The voices that are expressing concern about nanotechnology are the same voices that have expressed undue levels of concern about genetically modified organisms. As with nanoparticles, GMO's are neither inherently safe nor unsafe, and reasonable levels of regulation for safety are appropriate. However, none of the dire warnings about GMO's have come to pass. Already, African nations, such as Zambia and Zimbabwe, have rejected vitally needed food aid under pressure from European anti-GMO activists. The reflexive anti-technology stance that has been reflected in the GMO controversy will not be helpful in balancing the benefits and risks of nanoparticle technology.

In summary, I believe that existing regulatory mechanisms are sufficient to handle near-term applications of nanotechnology. As for the long term, we need to appreciate that a myriad of nanoscale technologies are inevitable. The current examinations and dialogues on achieving the promise while ameliorating the peril are appropriate and will deserve sharply increased attention as we get closer to realizing these revolutionary technologies.

Written Testimony

I am pleased to provide a more detailed written response to the issues raised by the committee. In this written portion of my response, I address the following issues:

- **Models of Technology Trends** : A discussion of why nanotechnology and related advanced technologies are inevitable. The underlying technologies are deeply integrated into our society and are advancing on many diverse fronts.
- **A Small Sample of Examples of True Nanotechnology**: a few of the implications of nanotechnology two to three decades from now.
- **The Economic Imperatives of the Law of Accelerating Returns** : the exponential advance of technology, including the accelerating miniaturization of technology, is driven by economic imperative, and, in turn, has a pervasive impact on the economy.
- **The Deeply Intertwined Promise and Peril of Nanotechnology and Related Advanced Technologies**: Technology is inherently a doubled-edged sword, and we will need to adopt strategies to encourage the benefits while ameliorating the risks. Relinquishing broad areas of technology, as has been proposed, is not

feasible and attempts to do so will only drive technology development underground, which will exacerbate the dangers.

MODELS OF TECHNOLOGY TRENDS	8
A SMALL SAMPLE OF EXAMPLES OF TRUE NANOTECHNOLOGY	24
THE ECONOMIC IMPERATIVES OF THE LAW OF ACCELERATING RETURNS	27
THE DEEPLY INTERTWINED PROMISE AND PERIL OF NANOTECHNOLOGY AND RELATED ADVANCED TECHNOLOGIES	36

Models of Technology Trends

A diverse technology such as nanotechnology progresses on many fronts and is comprised of hundreds of small steps forward, each benign in itself. An examination of these trends shows that technology in which the key features are measured in a small number of nanometers is inevitable. I hereby provide some examples of my study of technology trends.

The motivation for this study came from my interest in inventing. As an inventor in the 1970s, I came to realize that my inventions needed to make sense in terms of the enabling technologies and market forces that would exist when the invention was introduced, which would represent a very different world than when it was conceived. I began to develop models of how distinct technologies – electronics, communications, computer processors, memory, magnetic storage, and the size of technology – developed and how these changes rippled through markets and ultimately our social institutions. I realized that most inventions fail not because they never work, but because their timing is wrong. Inventing is a lot like surfing, you have to anticipate and catch the wave at just the right moment.

In the 1980s, my interest in technology trends and implications took on a life of its own, and I began to use my models of technology trends to project and anticipate the technologies of future times, such as the year 2000, 2010, 2020, and beyond. This enabled me to invent with the capabilities of the future. In the late 1980s, I wrote my first book, *The Age of Intelligent Machines*, which ended with the specter of machine intelligence becoming indistinguishable from its human progenitors. This book included hundreds of predictions about the 1990s and early 2000 years, and my track record of prediction has held up well.

During the 1990s I gathered empirical data on the apparent acceleration of all information-related technologies and sought to refine the mathematical models underlying these observations. In *The Age of Spiritual Machines* (ASM), which I wrote

in 1998, I introduced refined models of technology, and a theory I called “the law of accelerating returns,” which explained why technology evolves in an exponential fashion.

The Intuitive Linear View versus the Historical Exponential View

The future is widely misunderstood. Our forebears expected the future to be pretty much like their present, which had been pretty much like their past. Although exponential trends did exist a thousand years ago, they were at that very early stage where an exponential trend is so flat and so slow that it looks like no trend at all. So their lack of expectations was largely fulfilled. Today, in accordance with the common wisdom, everyone expects continuous technological progress and the social repercussions that follow. But the future will nonetheless be far more surprising than most observers realize because few have truly internalized the implications of the fact that the rate of change itself is accelerating.

Most long-range forecasts of technical feasibility in future time periods dramatically underestimate the power of future developments because they are based on what I call the “intuitive linear” view of history rather than the “historical exponential view.” To express this another way, it is not the case that we will experience a hundred years of progress in the twenty-first century; rather we will witness on the order of twenty thousand years of progress (at *today’s* rate of progress, that is).

When people think of a future period, they intuitively assume that the current rate of progress will continue for future periods. Even for those who have been around long enough to experience how the pace increases over time, an unexamined intuition nonetheless provides the impression that progress changes at the rate that we have experienced recently. From the mathematician’s perspective, a primary reason for this is that an exponential curve approximates a straight line when viewed for a brief duration. It is typical, therefore, that even sophisticated commentators, when considering the future, extrapolate the current pace of change over the next 10 years or 100 years to determine their expectations. This is why I call this way of looking at the future the “intuitive linear” view.

But a serious assessment of the history of technology shows that technological change is exponential. In exponential growth, we find that a key measurement such as computational power is multiplied by a constant factor for each unit of time (e.g., doubling every year) rather than just being added to incrementally. Exponential growth is a feature of any evolutionary process, of which technology is a primary example. One can examine the data in different ways, on different time scales, and for a wide variety of technologies ranging from electronic to biological, as well as social implications ranging from the size of the economy to human life span, and the acceleration of progress and growth applies. Indeed, we find not just simple exponential growth, but “double” exponential growth, meaning that the rate of exponential growth is itself growing exponentially. These observations do not rely merely on an assumption of the continuation of Moore’s law (i.e., the exponential shrinking of transistor sizes on an

integrated circuit), but is based on a rich model of diverse technological processes. What it clearly shows is that technology, particularly the pace of technological change, advances (at least) exponentially, not linearly, and has been doing so since the advent of technology, indeed since the advent of evolution on Earth.

Many scientists and engineers have what my colleague Lucas Hendrich calls “engineer’s pessimism.” Often an engineer or scientist who is so immersed in the difficulties and intricate details of a contemporary challenge fails to appreciate the ultimate long-term implications of their own work, and, in particular, the larger field of work that they operate in. Consider the biochemists in 1985 who were skeptical of the announcement of the goal of transcribing the entire genome in a mere 15 years. These scientists had just spent an entire year transcribing a mere one ten-thousandth of the genome, so even with reasonable anticipated advances, it seemed to them like it would be hundreds of years, if not longer, before the entire genome could be sequenced. Or consider the skepticism expressed in the mid 1980s that the Internet would ever be a significant phenomenon, given that it included only tens of thousands of nodes. The fact that the number of nodes was doubling every year and there were, therefore, likely to be tens of millions of nodes ten years later was not appreciated by those who struggled with “state of the art” technology in 1985, which permitted adding only a few thousand nodes throughout the world in a year.

I emphasize this point because it is the most important failure that would-be prognosticators make in considering future trends. The vast majority of technology forecasts and forecasters ignore altogether this “historical exponential view” of technological progress. Indeed, almost everyone I meet has a linear view of the future. That is why people tend to overestimate what can be achieved in the short term (because we tend to leave out necessary details), but underestimate what can be achieved in the long term (because the exponential growth is ignored).

The Law of Accelerating Returns

The ongoing acceleration of technology is the implication and inevitable result of what I call the “law of accelerating returns,” which describes the acceleration of the pace and the exponential growth of the products of an evolutionary process. This includes technology, particularly information-bearing technologies, such as computation. More specifically, the law of accelerating returns states the following:

✎ Evolution applies positive feedback in that the more capable methods resulting from one stage of evolutionary progress are used to create the next stage. As a result, the rate of progress of an evolutionary process increases exponentially over time. Over time, the “order” of the information embedded in the evolutionary process (i.e., the measure of how well the information fits a purpose, which in evolution is survival) increases.

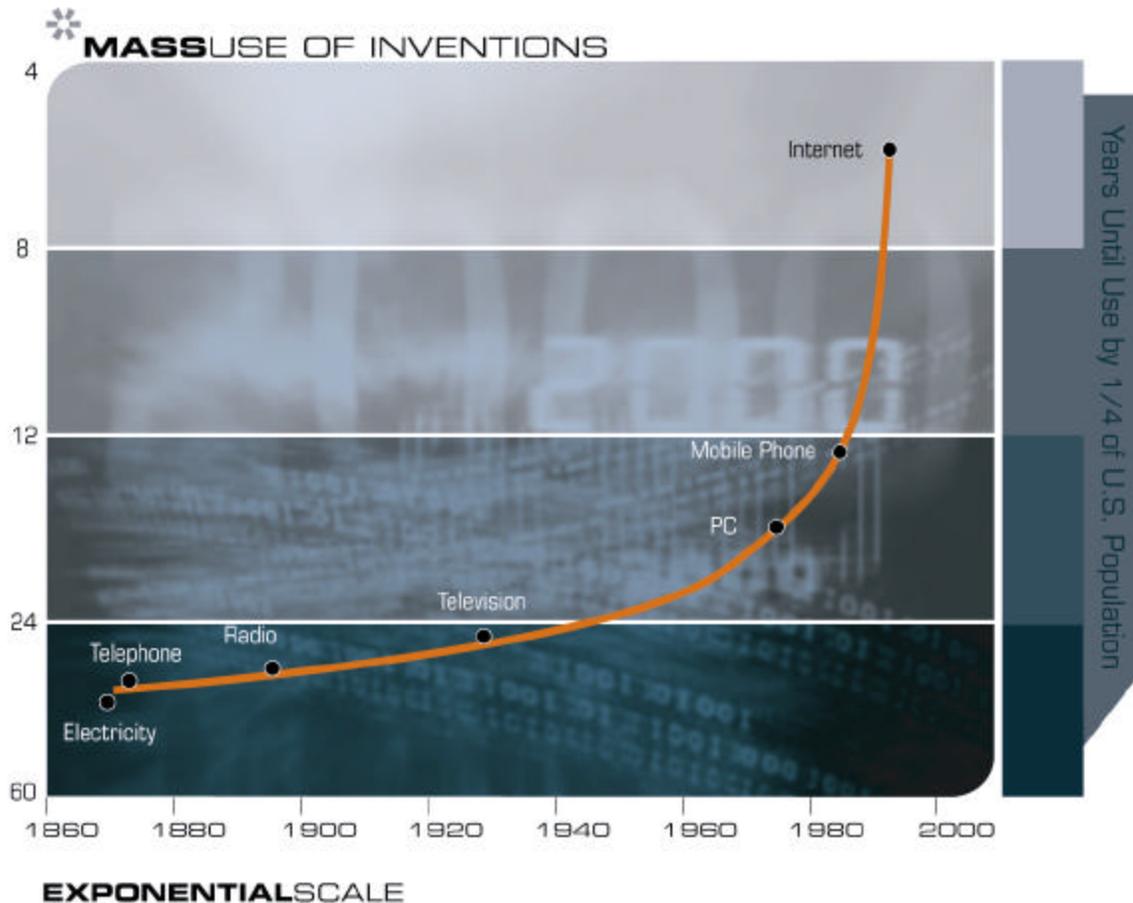
- ✧ A correlate of the above observation is that the “returns” of an evolutionary process (e.g., the speed, cost-effectiveness, or overall “power” of a process) increase exponentially over time.
- ✧ In another positive feedback loop, as a particular evolutionary process (e.g., computation) becomes more effective (e.g., cost effective), greater resources are deployed towards the further progress of that process. This results in a second level of exponential growth (i.e., the rate of exponential growth itself grows exponentially).
- ✧ Biological evolution is one such evolutionary process.
- ✧ Technological evolution is another such evolutionary process. Indeed, the emergence of the first technology-creating species resulted in the new evolutionary process of technology. Therefore, technological evolution is an outgrowth of – and a continuation of – biological evolution.
- ✧ A specific paradigm (a method or approach to solving a problem, e.g., shrinking transistors on an integrated circuit as an approach to making more powerful computers) provides exponential growth until the method exhausts its potential. When this happens, a paradigm shift (a fundamental change in the approach) occurs, which enables exponential growth to continue.
- ✧ Each paradigm follows an “S-curve,” which consists of slow growth (the early phase of exponential growth), followed by rapid growth (the late, explosive phase of exponential growth), followed by a leveling off as the particular paradigm matures.
- ✧ During this third or maturing phase in the life cycle of a paradigm, pressure builds for the next paradigm shift.
- ✧ When the paradigm shift occurs, the process begins a new S-curve.
- ✧ Thus the acceleration of the overall evolutionary process proceeds as a sequence of S-curves, and the overall exponential growth consists of this cascade of S-curves.
- ✧ The resources underlying the exponential growth of an evolutionary process are relatively unbounded.
- ✧ One resource is the (ever-growing) order of the evolutionary process itself. Each stage of evolution provides more powerful tools for the next. In biological evolution, the advent of DNA allowed more powerful and faster evolutionary “experiments.” Later, setting the “designs” of animal body plans during the Cambrian explosion allowed rapid evolutionary development of other body organs, such as the brain. Or to take a more recent example, the advent of computer-assisted design tools allows rapid development of the next generation of computers.

✂ The other required resource is the “chaos” of the environment in which the evolutionary process takes place and which provides the options for further diversity. In biological evolution, diversity enters the process in the form of mutations and ever-changing environmental conditions, including cosmological disasters (e.g., asteroids hitting the Earth). In technological evolution, human ingenuity combined with ever-changing market conditions keep the process of innovation going.

If we apply these principles at the highest level of evolution on Earth, the first step, the creation of cells, introduced the paradigm of biology. The subsequent emergence of DNA provided a digital method to record the results of evolutionary experiments. Then, the evolution of a species that combined rational thought with an opposable appendage (the thumb) caused a fundamental paradigm shift from biology to technology. The upcoming primary paradigm shift will be from biological thinking to a hybrid combining biological and nonbiological thinking. This hybrid will include “biologically inspired” processes resulting from the reverse engineering of biological brains.

If we examine the timing of these steps, we see that the process has continuously accelerated. The evolution of life forms required billions of years for the first steps (e.g., primitive cells); later on progress accelerated. During the Cambrian explosion, major paradigm shifts took only tens of millions of years. Later on, Humanoids developed over a period of millions of years, and Homo sapiens over a period of only hundreds of thousands of years.

With the advent of a technology-creating species, the exponential pace became too fast for evolution through DNA-guided protein synthesis and moved on to human-created technology. Technology goes beyond mere tool making; it is a process of creating ever more powerful technology using the tools from the previous round of innovation, and is, thereby, an evolutionary process. The first technological steps -- sharp edges, fire, the wheel -- took tens of thousands of years. For people living in this era, there was little noticeable technological change in even a thousand years. By 1000 AD, progress was much faster and a paradigm shift required only a century or two. In the nineteenth century, we saw more technological change than in the nine centuries preceding it. Then in the first twenty years of the twentieth century, we saw more advancement than in all of the nineteenth century. Now, paradigm shifts occur in only a few years time. The World Wide Web did not exist in anything like its present form just a few years ago; it didn't exist at all a decade ago.



The paradigm shift rate (i.e., the overall rate of technical progress) is currently doubling (approximately) every decade; that is, paradigm shift times are halving every decade (and the rate of acceleration is itself growing exponentially). So, the technological progress in the twenty-first century will be equivalent to what would require (in the linear view) on the order of 200 centuries. In contrast, the twentieth century saw only about 20 years of progress (again at today's rate of progress) since we have been speeding up to current rates. So the twenty-first century will see about a thousand times greater technological change than its predecessor.

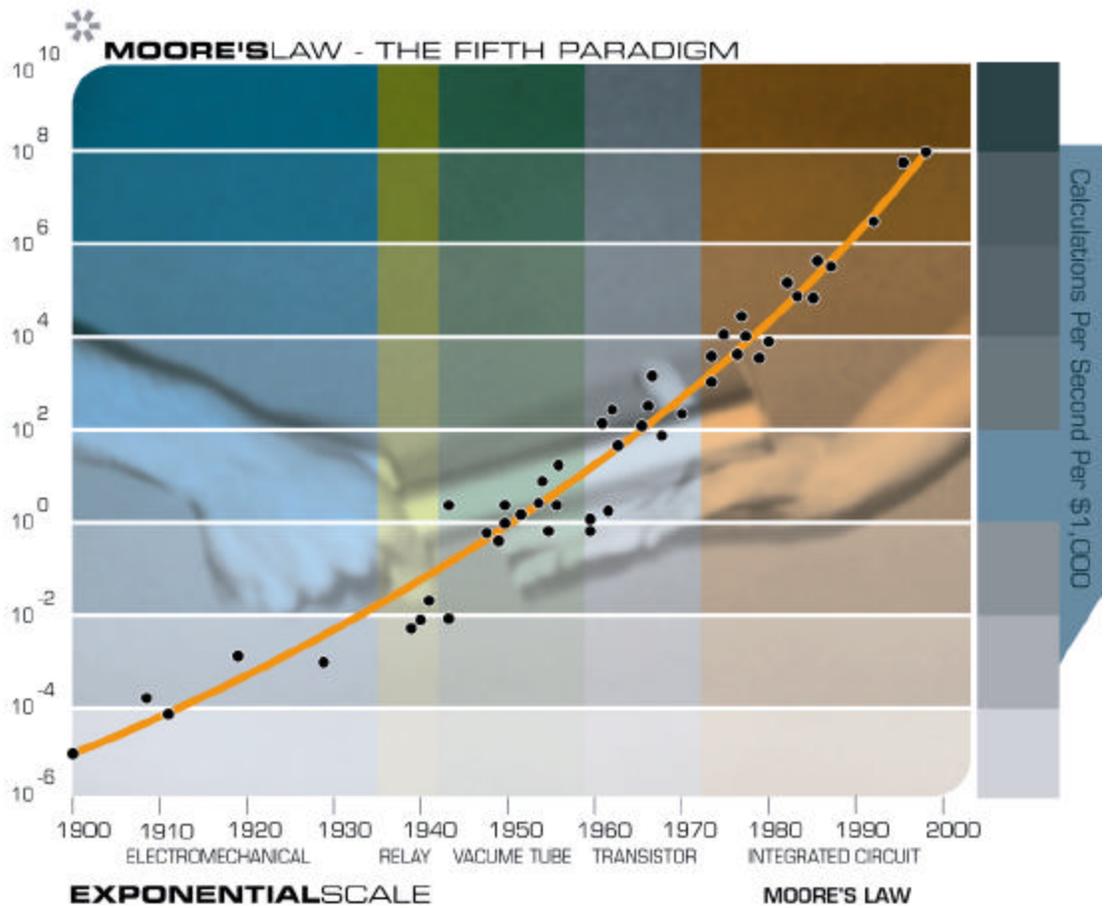
Moore's Law and Beyond

There is a wide range of technologies that are subject to the law of accelerating returns. The exponential trend that has gained the greatest public recognition has become known as "Moore's Law." Gordon Moore, one of the inventors of integrated circuits, and then Chairman of Intel, noted in the mid-1970s that we could squeeze twice as many transistors on an integrated circuit every 24 months. Given that the electrons have less distance to travel, the circuits also run twice as fast, providing an overall quadrupling of computational power.

However, the exponential growth of computing is much broader than Moore's Law.

If we plot the speed (in instructions per second) per \$1000 (in constant dollars) of 49 famous calculators and computers spanning the entire twentieth century, we note that there were four completely different paradigms that provided exponential growth in the price-performance of computing before the integrated circuits were invented. Therefore, Moore's Law was not the first, but the fifth paradigm to exponentially grow the power of computation. And it won't be the last. When Moore's Law reaches the end of its S-Curve, now expected before 2020, the exponential growth will continue with three-dimensional molecular computing, a prime example of the application of nanotechnology, which will constitute the sixth paradigm.

When I suggested in my book *The Age of Spiritual Machines*, published in 1999, that three-dimensional molecular computing, particularly an approach based on using carbon nanotubes, would become the dominant computing hardware technology in the teen years of this century, that was considered a radical notion. There has been so much progress in the past four years, with literally dozens of major milestones having been achieved, that this expectation is now a mainstream view.



*Moore's Law Was Not the First, but the Fifth Paradigm
To Provide Exponential Growth of Computing
Each time one paradigm runs out of steam, another picks up the pace*

The exponential growth of computing is a marvelous quantitative example of the exponentially growing returns from an evolutionary process. We can express the exponential growth of computing in terms of an accelerating pace: it took 90 years to achieve the first MIPS (million instructions per second) per thousand dollars; now we add one MIPS per thousand dollars every day.

Moore's Law narrowly refers to the number of transistors on an integrated circuit of fixed size, and sometimes has been expressed even more narrowly in terms of transistor feature size. But rather than feature size (which is only one contributing factor), or even number of transistors, I think the most appropriate measure to track is computational speed per unit cost. This takes into account many levels of "cleverness" (i.e., innovation, which is to say, technological evolution). In addition to all of the innovation in integrated circuits, there are multiple layers of innovation in computer

design, e.g., pipelining, parallel processing, instruction look-ahead, instruction and memory caching, and many others.

The human brain uses a very inefficient electrochemical digital-controlled analog computational process. The bulk of the calculations are done in the interneuronal connections at a speed of only about 200 calculations per second (in each connection), which is about ten million times slower than contemporary electronic circuits. But the brain gains its prodigious powers from its extremely parallel organization *in three dimensions*. There are many technologies in the wings that build circuitry in three dimensions. Nanotubes, an example of nanotechnology, which is already working in laboratories, build circuits from pentagonal arrays of carbon atoms. One cubic inch of nanotube circuitry would be a million times more powerful than the human brain. There are more than enough new computing technologies now being researched, including three-dimensional silicon chips, optical and silicon spin computing, crystalline computing, DNA computing, and quantum computing, to keep the law of accelerating returns as applied to computation going for a long time.

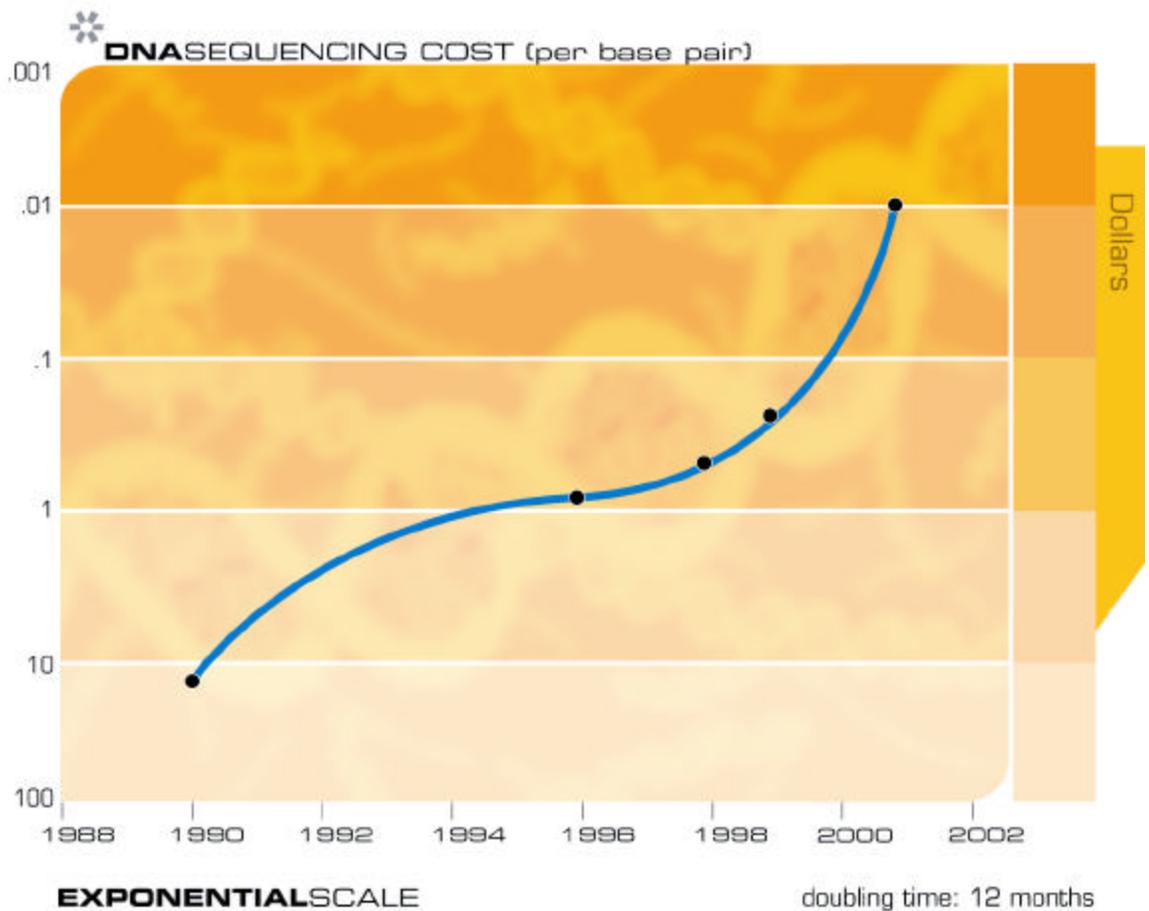
As I discussed above, it is important to distinguish between the “S” curve (an “S” stretched to the right, comprising very slow, virtually unnoticeable growth – followed by very rapid growth – followed by a flattening out as the process approaches an asymptote) that is characteristic of any specific technological paradigm and the continuing exponential growth that is characteristic of the ongoing evolutionary process of technology. Specific paradigms, such as Moore’s Law, do ultimately reach levels at which exponential growth is no longer feasible. That is why Moore’s Law is an S curve. But the growth of computation is an ongoing exponential (at least until we “saturate” the Universe with the intelligence of our human-machine civilization, but that will not be a limit in this coming century). In accordance with the law of accelerating returns, paradigm shift, also called innovation, turns the S curve of any specific paradigm into a continuing exponential. A new paradigm (e.g., three-dimensional circuits) takes over when the old paradigm approaches its natural limit, which has already happened at least four times in the history of computation. This difference also distinguishes the tool making of non-human species, in which the mastery of a tool-making (or using) skill by each animal is characterized by an abruptly ending S shaped learning curve, versus human-created technology, which has followed an exponential pattern of growth and acceleration since its inception.

DNA Sequencing, Memory, Communications, the Internet, and Miniaturization

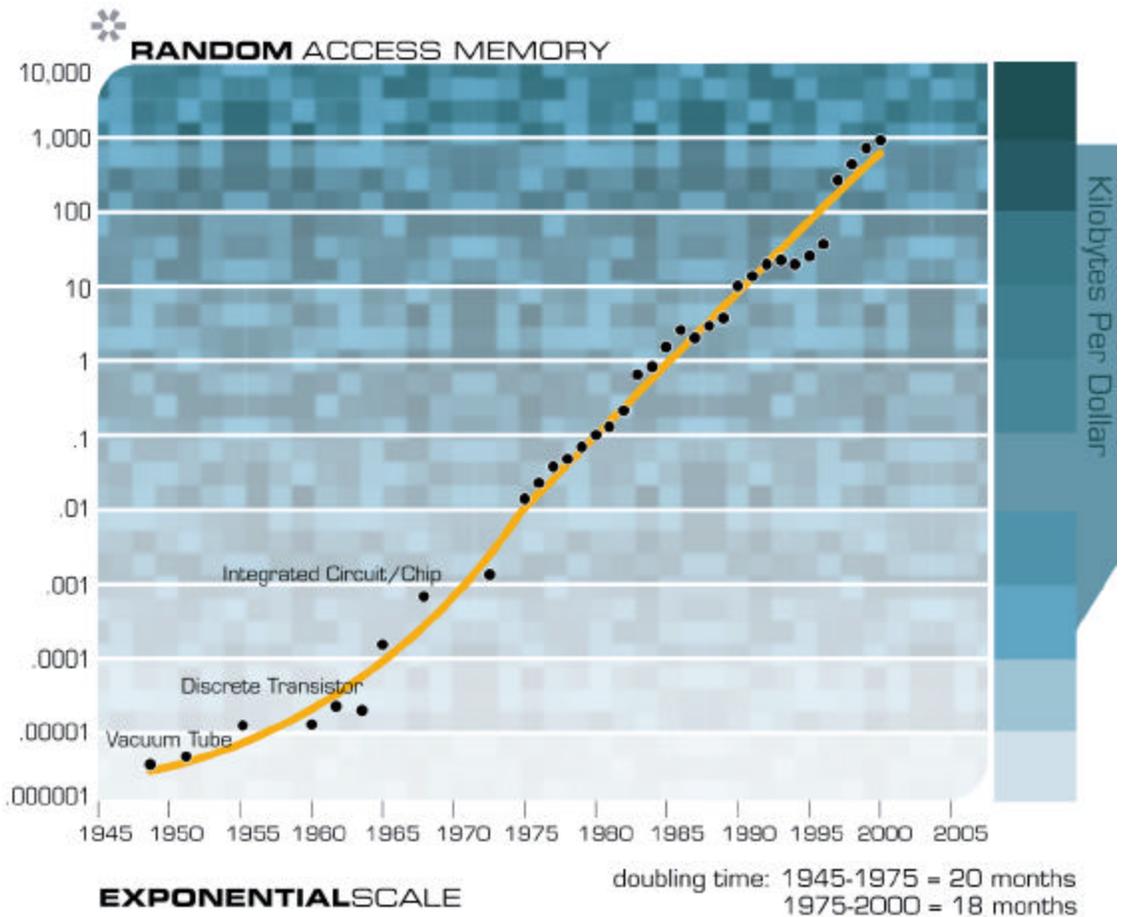
This “law of accelerating returns” applies to all of technology, indeed to any true evolutionary process, and can be measured with remarkable precision in information-based technologies. There are a great many examples of the exponential growth implied by the law of accelerating returns in technologies, as varied as DNA sequencing, communication speeds, brain scanning, electronics of all kinds, and even in the rapidly shrinking size of technology, which is directly relevant to the discussion at this hearing.

The future nanotechnology age results not from the exponential explosion of computation alone, but rather from the interplay and myriad synergies that will result from manifold intertwined technological revolutions. Also, keep in mind that every point on the exponential growth curves underlying these panoply of technologies (see the graphs below) represents an intense human drama of innovation and competition. It is remarkable therefore that these chaotic processes result in such smooth and predictable exponential trends.

As I noted above, when the human genome scan started fourteen years ago, critics pointed out that given the speed with which the genome could then be scanned, it would take thousands of years to finish the project. Yet the fifteen year project was nonetheless completed slightly ahead of schedule.

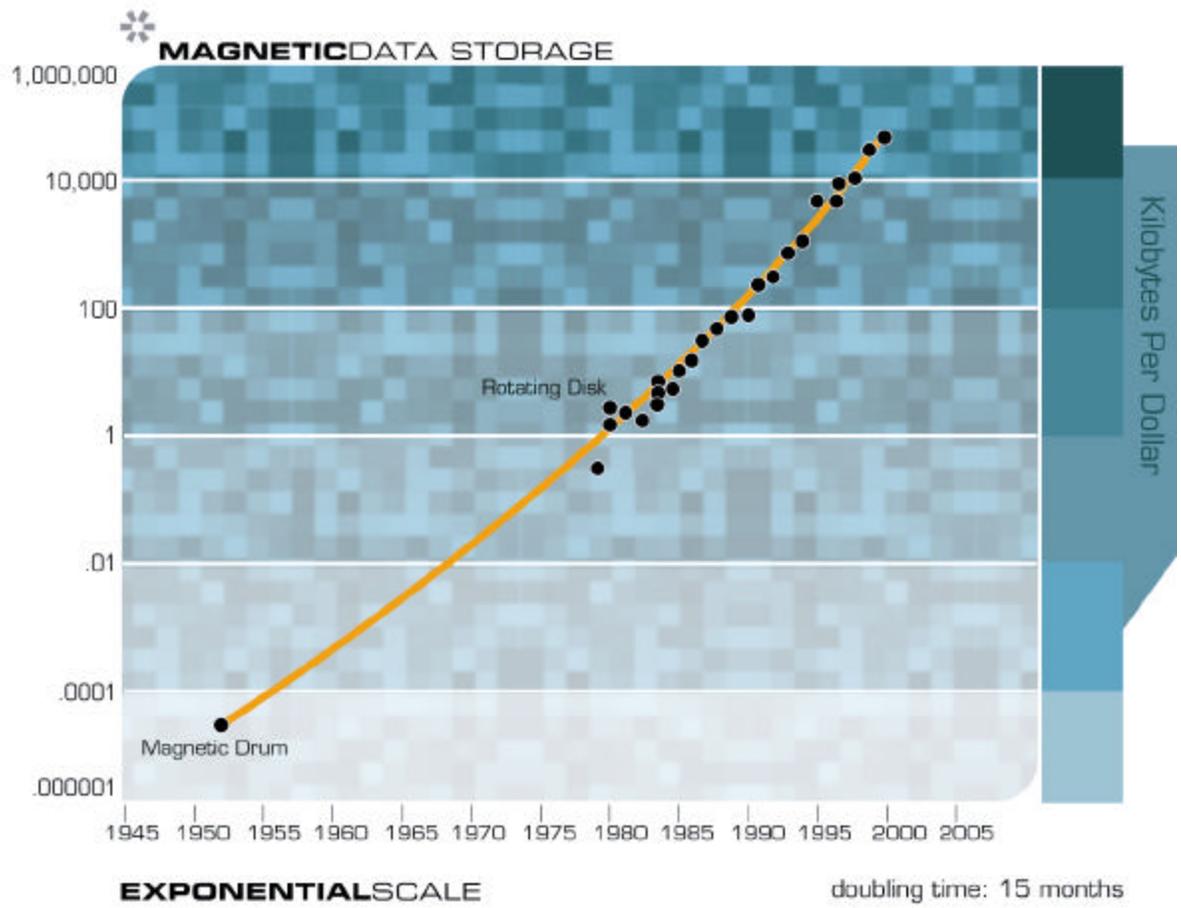


Of course, we expect to see exponential growth in electronic memories such as RAM.

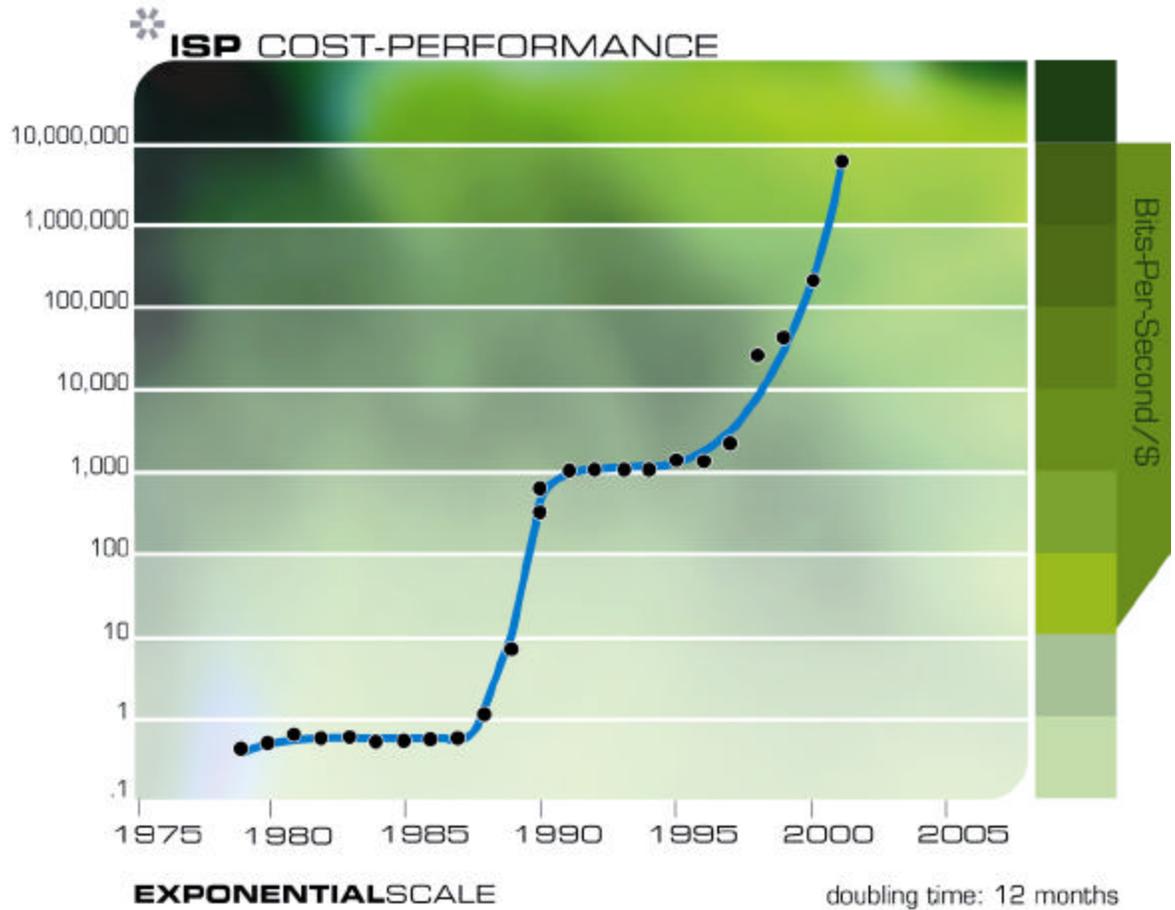


Notice How Exponential Growth Continued through Paradigm Shifts from Vacuum Tubes to Discrete Transistors to Integrated Circuits

However, growth in magnetic memory is not primarily a matter of Moore's law, but includes advances in mechanical and electromagnetic systems.



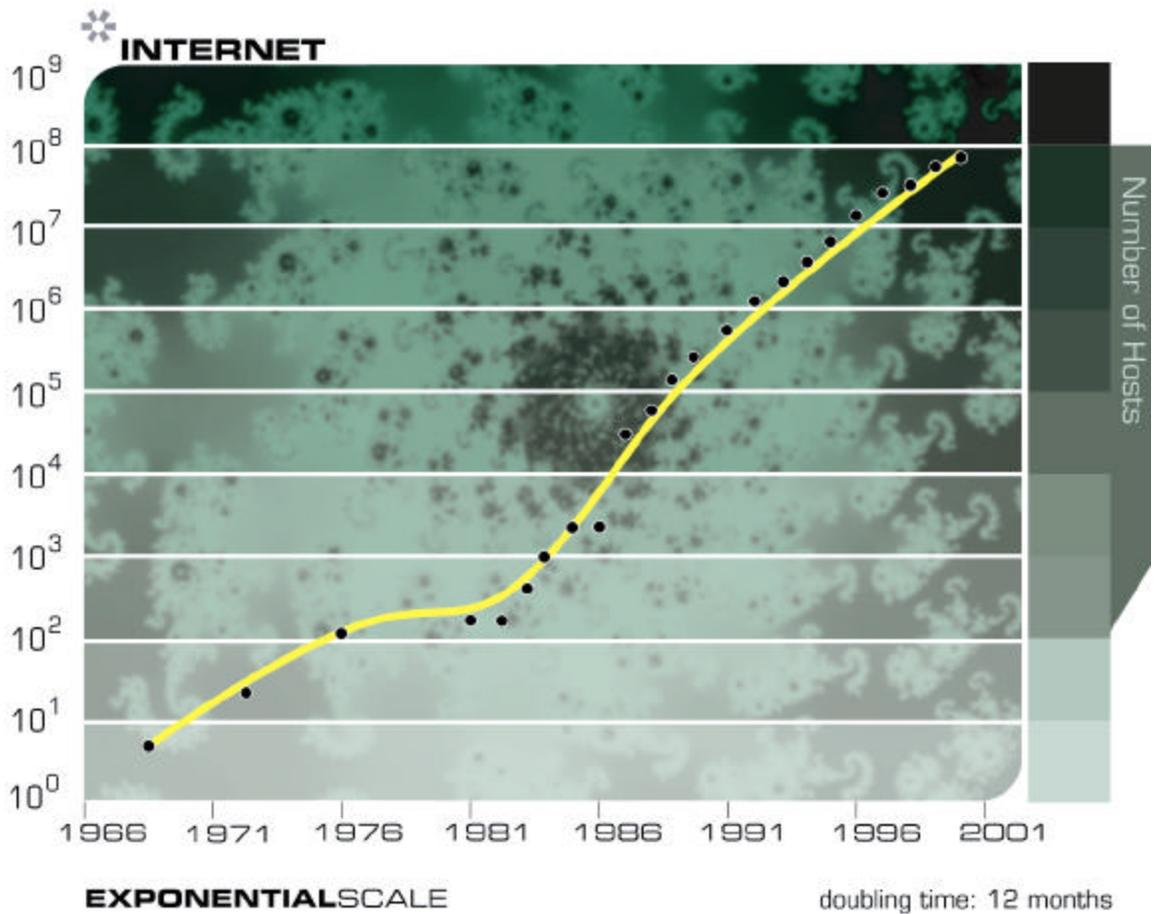
Exponential growth in communications technology has been even more explosive than in computation and is no less significant in its implications. Again, this progression involves far more than just shrinking transistors on an integrated circuit, but includes accelerating advances in fiber optics, optical switching, electromagnetic technologies, and others.

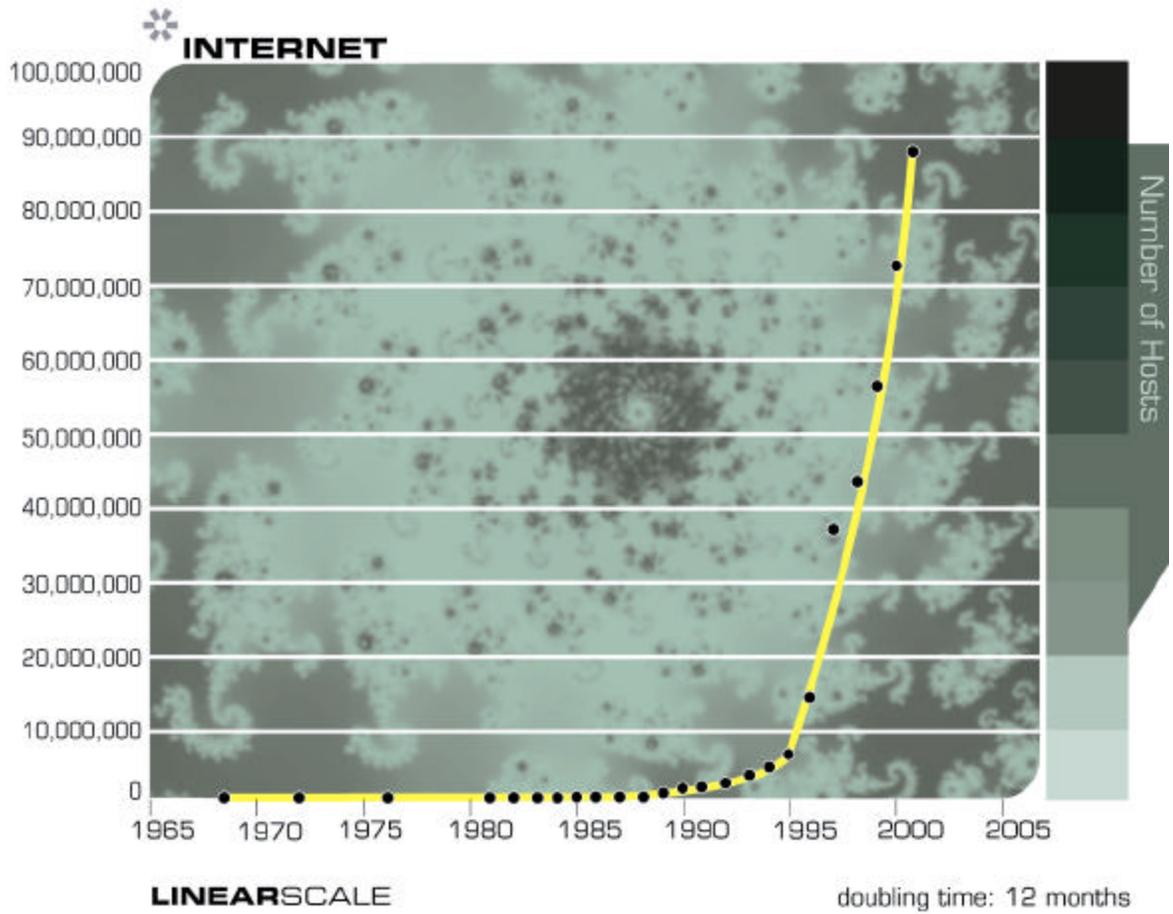


Notice Cascade of “S” Curves

Note that in the above chart we can actually see the progression of “S” curves: the acceleration fostered by a new paradigm, followed by a leveling off as the paradigm runs out of steam, followed by renewed acceleration through paradigm shift.

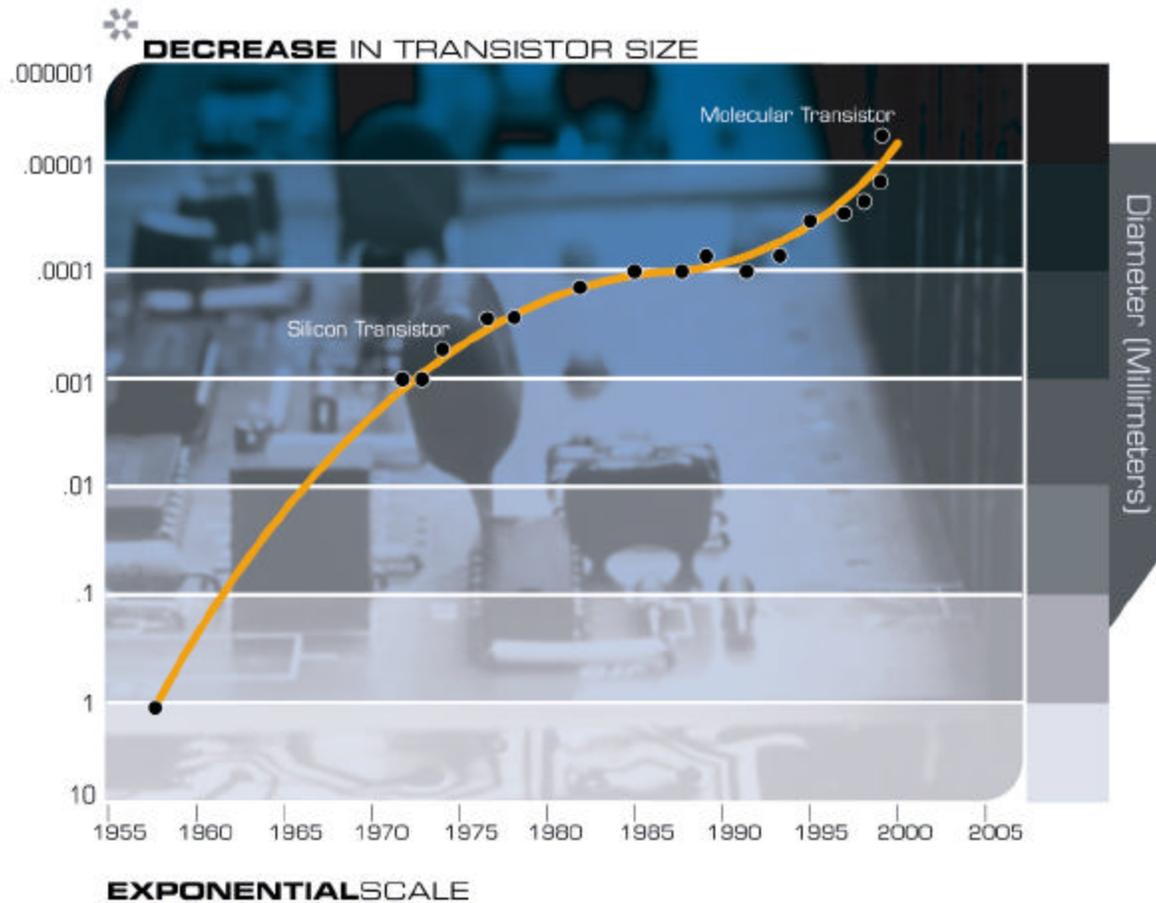
The following two charts show the overall growth of the Internet based on the number of hosts (server computers). These two charts plot the same data, but one is on an exponential axis and the other is linear. As I pointed out earlier, whereas technology progresses in the exponential domain, we experience it in the linear domain. So from the perspective of most observers, nothing was happening until the mid 1990s when seemingly out of nowhere, the World Wide Web and email exploded into view. But the emergence of the Internet into a worldwide phenomenon was readily predictable much earlier by examining the exponential trend data.

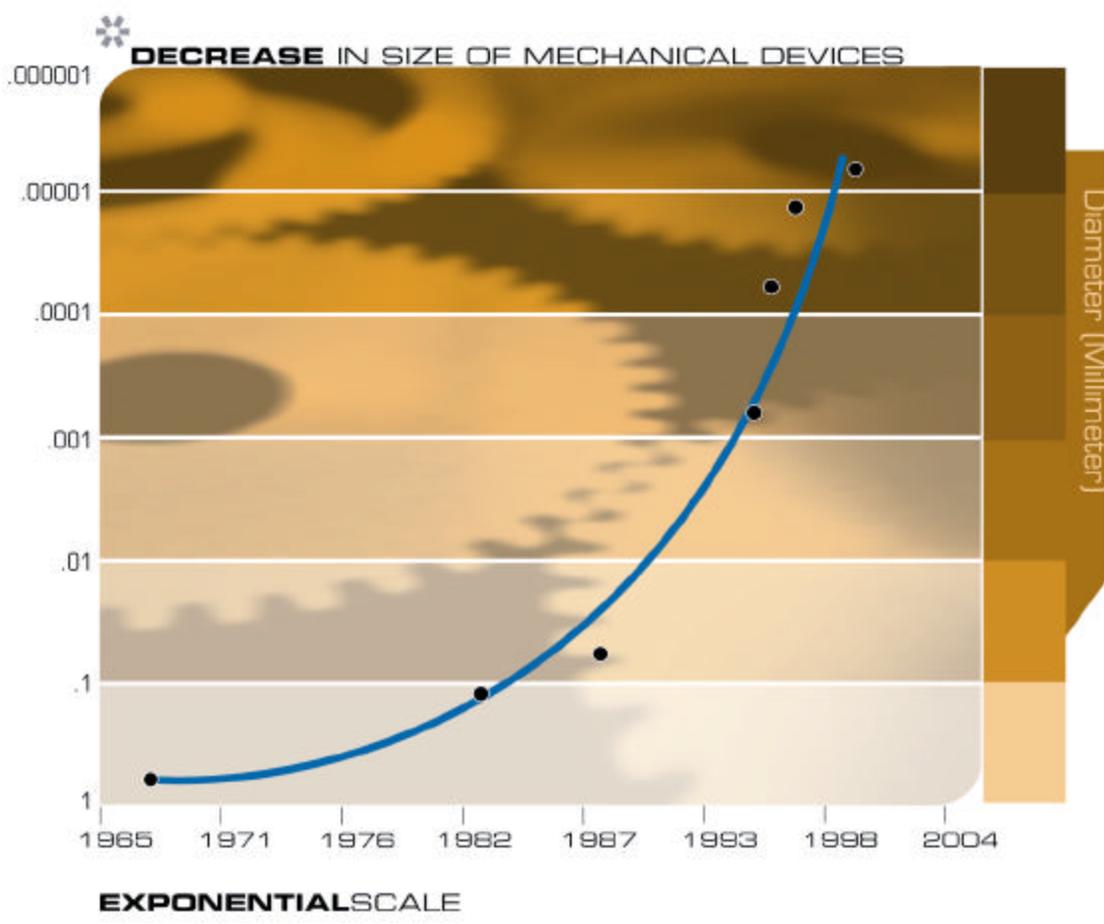




Notice how the explosion of the Internet appears to be a surprise from the Linear Chart, but was perfectly predictable from the Exponential Chart

The most relevant trend to this hearing, and one that will have profound implications for the twenty-first century is the pervasive trend towards making things smaller, i.e., miniaturization. The salient implementation sizes of a broad range of technologies, both electronic and mechanical, are shrinking, also at a double-exponential rate. At present, we are shrinking technology by a factor of approximately 5.6 per linear dimension per decade.





A Small Sample of Examples of True Nanotechnology

Ubiquitous nanotechnology is two to three decades away. A prime example of its application will be to deploy billions of “nanobots”: small robots the size of human blood cells that can travel inside the human bloodstream. This notion is not as futuristic as it may sound in that there have already been successful animal experiments using this concept . There are already four major conferences on “BioMEMS” (Biological Micro Electronic Mechanical Systems) covering devices in the human blood stream.

Consider several examples of nanobot technology, which, based on miniaturization and cost reduction trends, will be feasible within 30 years. In addition to scanning the human brain to facilitate human brain reverse engineering, these nanobots will be able to perform a broad variety of diagnostic and therapeutic functions inside the

bloodstream and human body. Robert Freitas, for example, has designed robotic replacements for human blood cells that perform hundreds or thousands of times more effectively than their biological counterparts. With Freitas' "respirocytes," (robotic red blood cells), you could do an Olympic sprint for 15 minutes without taking a breath. His robotic macrophages will be far more effective than our white blood cells at combating pathogens. His DNA repair robot would be able to repair DNA transcription errors, and even implement needed DNA changes. Although Freitas' conceptual designs are two or three decades away, there has already been substantial progress on bloodstream-based devices. For example, one scientist has cured type I Diabetes in rats with a nanoengineered device that incorporates pancreatic Islet cells. The device has seven-nanometer pores that let insulin out, but block the antibodies which destroy these cells. There are many innovative projects of this type already under way.

Clearly, nanobot technology has profound military applications, and any expectation that such uses will be "relinquished" are highly unrealistic. Already, DOD is developing "smart dust," which are tiny robots the size of insects or even smaller. Although not quite nanotechnology, millions of these devices can be dropped into enemy territory to provide highly detailed surveillance. The potential application for even smaller, nanotechnology-based devices is even greater. Want to find Saddam Hussein or Osama bin Laden? Need to locate hidden weapons of mass destruction? Billions of essentially invisible spies could monitor every square inch of enemy territory, identify every person and every weapon, and even carry out missions to destroy enemy targets. The only way for an enemy to counteract such a force is, of course, with their own nanotechnology. The point is that nanotechnology-based weapons will obsolete weapons of larger size.

In addition, nanobots will also be able to expand our experiences and our capabilities. Nanobot technology will provide fully immersive, totally convincing virtual reality in the following way. The nanobots take up positions in close physical proximity to every interneuronal connection coming from all of our senses (e.g., eyes, ears, skin). We already have the technology for electronic devices to communicate with neurons in both directions that requires no direct physical contact with the neurons. For example, scientists at the Max Planck Institute have developed "neuron transistors" that can detect the firing of a nearby neuron, or alternatively, can cause a nearby neuron to fire, or suppress it from firing. This amounts to two-way communication between neurons and the electronic-based neuron transistors. The Institute scientists demonstrated their invention by controlling the movement of a living leech from their computer. Again, the primary aspect of nanobot-based virtual reality that is not yet feasible is size and cost.

When we want to experience real reality, the nanobots just stay in position (in the capillaries) and do nothing. If we want to enter virtual reality, they suppress all of the inputs coming from the real senses, and replace them with the signals that would be appropriate for the virtual environment. You (i.e., your brain) could decide to cause your muscles and limbs to move as you normally would, but the nanobots again intercept these interneuronal signals, suppress your real limbs from moving, and instead cause your

virtual limbs to move and provide the appropriate movement and reorientation in the virtual environment.

The Web will provide a panoply of virtual environments to explore. Some will be recreations of real places, others will be fanciful environments that have no “real” counterpart. Some indeed would be impossible in the physical world (perhaps, because they violate the laws of physics). We will be able to “go” to these virtual environments by ourselves, or we will meet other people there, both real people and simulated people. Of course, ultimately there won’t be a clear distinction between the two.

By 2030, going to a web site will mean entering a full-immersion virtual-reality environment. In addition to encompassing all of the senses, these shared environments can include emotional overlays as the nanobots will be capable of triggering the neurological correlates of emotions, sexual pleasure, and other derivatives of our sensory experience and mental reactions.

In the same way that people today beam their lives from web cams in their bedrooms, “experience beamers” circa 2030 will beam their entire flow of sensory experiences, and if so desired, their emotions and other secondary reactions. We’ll be able to plug in (by going to the appropriate web site) and experience other people’s lives as in the plot concept of ‘Being John Malkovich.’ Particularly interesting experiences can be archived and relived at any time.

We won’t need to wait until 2030 to experience shared virtual-reality environments, at least for the visual and auditory senses. Full-immersion visual-auditory environments will be available by the end of this decade, with images written directly onto our retinas by our eyeglasses and contact lenses. All of the electronics for the computation, image reconstruction, and very high bandwidth wireless connection to the Internet will be embedded in our glasses and woven into our clothing, so computers as distinct objects will disappear.

In my view, the most significant implication of the development of nanotechnology and related advanced technologies of the 21st century will be the merger of biological and nonbiological intelligence. First, it is important to point out that well before the end of the twenty-first century, thinking on nonbiological substrates will dominate. Biological thinking is stuck at 10^{26} calculations per second (for all biological human brains), and that figure will not appreciably change, even with bioengineering changes to our genome. Nonbiological intelligence, on the other hand, is growing at a double-exponential rate and will vastly exceed biological intelligence well before the middle of this century. However, in my view, this nonbiological intelligence should still be considered human as it is fully derivative of the human-machine civilization. The merger of these two worlds of intelligence is not merely a merger of biological and nonbiological thinking mediums, but more importantly one of method and organization of thinking.

One of the key ways in which the two worlds can interact will be through nanobots. Nanobot technology will be able to expand our minds in virtually any imaginable way. Our brains today are relatively fixed in design. Although we do add patterns of interneuronal connections and neurotransmitter concentrations as a normal part of the learning process, the current overall capacity of the human brain is highly constrained, restricted to a mere hundred trillion connections. Brain implants based on massively distributed intelligent nanobots will ultimately expand our memories a trillion fold, and otherwise vastly improve all of our sensory, pattern recognition, and cognitive abilities. Since the nanobots are communicating with each other over a wireless local area network, they can create any set of new neural connections, can break existing connections (by suppressing neural firing), can create new hybrid biological-nonbiological networks, as well as add vast new nonbiological networks.

Using nanobots as brain extenders is a significant improvement over the idea of surgically installed neural implants, which are beginning to be used today (e.g., ventral posterior nucleus, subthalamic nucleus, and ventral lateral thalamus neural implants to counteract Parkinson's Disease and tremors from other neurological disorders, cochlear implants, and others.) Nanobots will be introduced without surgery, essentially just by injecting or even swallowing them. They can all be directed to leave, so the process is easily reversible. They are programmable, in that they can provide virtual reality one minute, and a variety of brain extensions the next. They can change their configuration, and clearly can alter their software. Perhaps most importantly, they are massively distributed and therefore can take up billions or trillions of positions throughout the brain, whereas a surgically introduced neural implant can only be placed in one or at most a few locations.

The Economic Imperatives of the Law of Accelerating Returns

It is the economic imperative of a competitive marketplace that is driving technology forward and fueling the law of accelerating returns. In turn, the law of accelerating returns is transforming economic relationships.

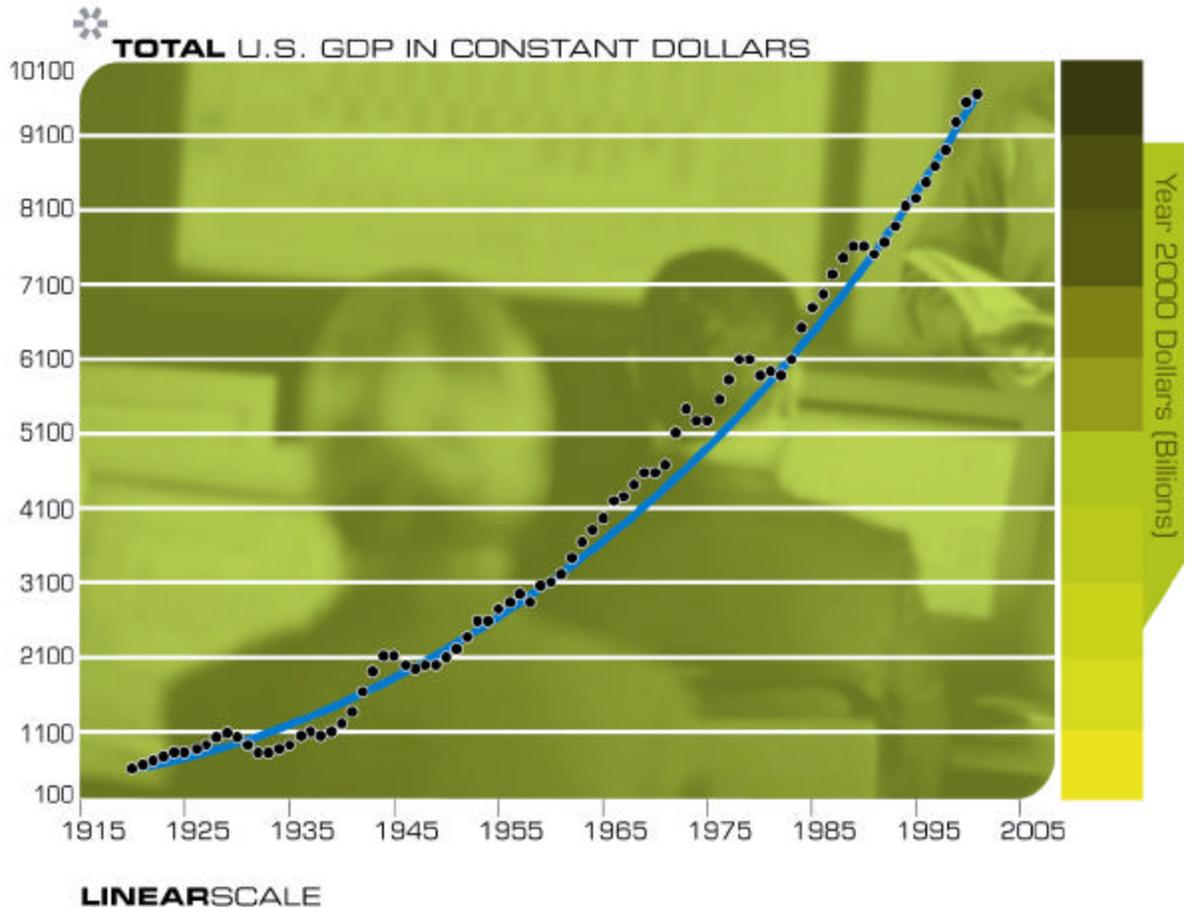
The primary force driving technology is economic imperative. We are moving towards nanoscale machines, as well as more intelligent machines, as the result of a myriad of small advances, each with their own particular economic justification.

To use one small example of many from my own experience at one of my companies (Kurzweil Applied Intelligence), whenever we came up with a slightly more intelligent version of speech recognition, the new version invariably had greater value than the earlier generation and, as a result, sales increased. It is interesting to note that in the example of speech recognition software, the three primary surviving competitors stayed very close to each other in the intelligence of their software. A few other companies that failed to do so (e.g., Speech Systems) went out of business. At any point

in time, we would be able to sell the version prior to the latest version for perhaps a quarter of the price of the current version. As for versions of our technology that were two generations old, we couldn't even give those away.

There is a vital economic imperative to create smaller and more intelligent technology. Machines that can more precisely carry out their missions have enormous value. That is why they are being built. There are tens of thousands of projects that are advancing the various aspects of the law of accelerating returns in diverse incremental ways. Regardless of near-term business cycles, the support for "high tech" in the business community, and in particular for software advancement, has grown enormously. When I started my optical character recognition (OCR) and speech synthesis company (Kurzweil Computer Products, Inc.) in 1974, high-tech venture deals totaled approximately \$10 million. Even during today's high tech recession, the figure is 100 times greater. We would have to repeal capitalism and every visage of economic competition to stop this progression.

The economy (viewed either in total or per capita) has been growing exponentially throughout this century:



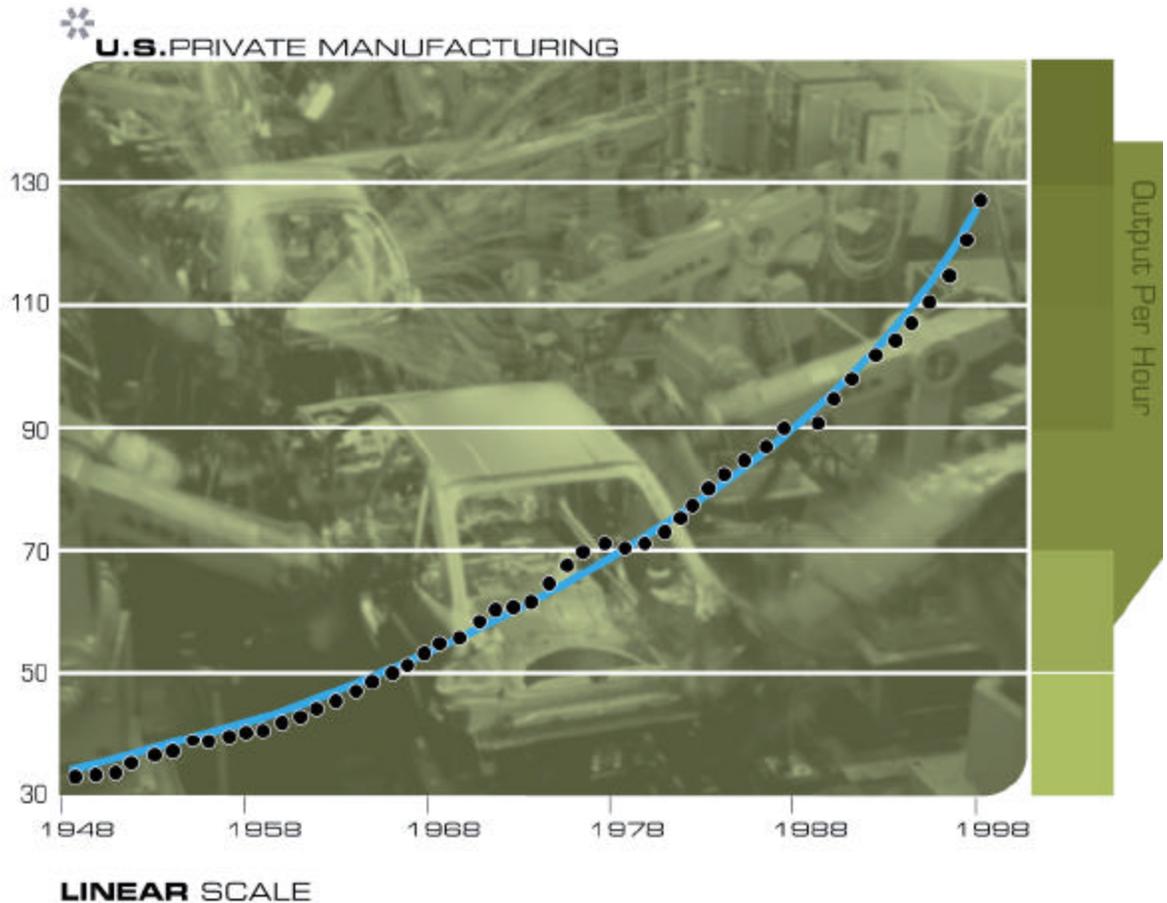
Note that the underlying exponential growth in the economy is a far more powerful force than periodic recessions. Even the “Great Depression” represents only a minor blip compared to the underlying pattern of growth. Most importantly, recessions, including the depression, represent only temporary deviations from the underlying curve. In each case, the economy ends up exactly where it would have been had the recession/depression never occurred.

Productivity (economic output per worker) has also been growing exponentially. Even these statistics are greatly understated because they do not fully reflect significant improvements in the quality and features of products and services. It is not the case that “a car is a car;” there have been significant improvements in safety, reliability, and features. Certainly, \$1000 of computation today is immeasurably more powerful than \$1000 of computation ten years ago (by a factor of more than 1000). There are a myriad of such examples. Pharmaceutical drugs are increasingly effective. Products ordered in five minutes on the web and delivered to your door are worth more than products that you have to fetch yourself. Clothes custom-manufactured for your unique body scan are worth more than clothes you happen to find left on a store rack. These sorts of

improvements are true for most product categories, and none of them are reflected in the productivity statistics.

The statistical methods underlying the productivity measurements tend to factor out gains by essentially concluding that we still only get one dollar of products and services for a dollar despite the fact that we get much more for a dollar (e.g., compare a \$1,000 computer today to one ten years ago). University of Chicago Professor Pete Klenow and University of Rochester Professor Mark Bilts estimate that the value of existing goods has been increasing at 1.5% per year for the past 20 years because of qualitative improvements. This still does not account for the introduction of entirely new products and product categories (e.g., cell phones, pagers, pocket computers). The Bureau of Labor Statistics, which is responsible for the inflation statistics, uses a model that incorporates an estimate of quality growth at only 0.5% per year, reflecting a systematic underestimate of quality improvement and a resulting overestimate of inflation by at least 1 percent per year.

Despite these weaknesses in the productivity statistical methods, the gains in productivity are now reaching the steep part of the exponential curve. Labor productivity grew at 1.6% per year until 1994, then rose at 2.4% per year, and is now growing even more rapidly. In the quarter ending July 30, 2000, labor productivity grew at 5.3%. Manufacturing productivity grew at 4.4% annually from 1995 to 1999, durables manufacturing at 6.5% per year.



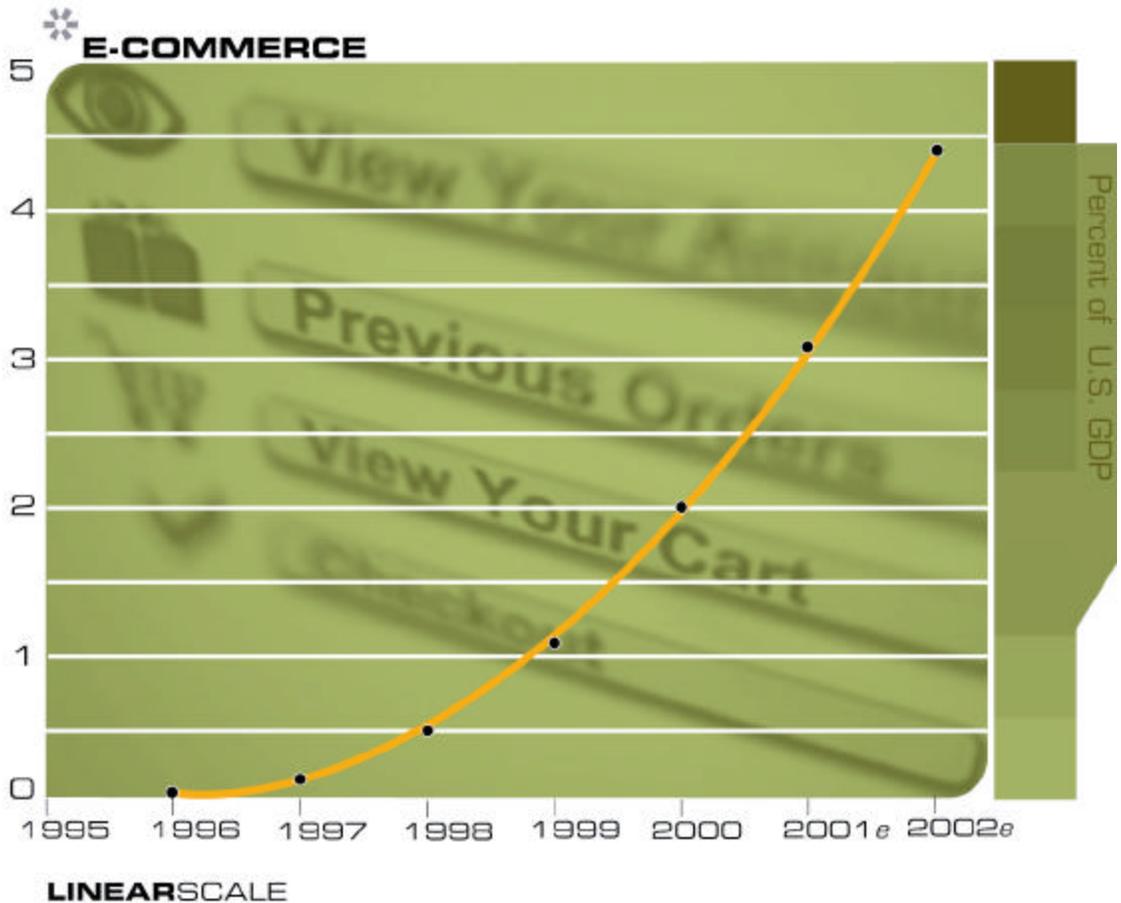
The 1990s have seen the most powerful deflationary forces in history. This is why we are not seeing inflation. Yes, it's true that low unemployment, high asset values, economic growth, and other such factors are inflationary, but these factors are offset by the double-exponential trends in the price-performance of all information-based technologies: computation, memory, communications, biotechnology, miniaturization, and even the overall rate of technical progress. These technologies deeply affect all industries. We are also undergoing massive disintermediation in the channels of distribution through the Web and other new communication technologies, as well as escalating efficiencies in operations and administration.

All of the technology trend charts above represent massive deflation. There are many examples of the impact of these escalating efficiencies. BP Amoco's cost for finding oil is now less than \$1 per barrel, down from nearly \$10 in 1991. Processing an Internet transaction costs a bank one penny, compared to over \$1 using a teller ten years ago. A Roland Berger/Deutsche Bank study estimates a cost savings of \$1200 per North American car over the next five years. A more optimistic Morgan Stanley study estimates that Internet-based procurement will save Ford, GM, and DaimlerChrysler about \$2700 per vehicle.

It is important to point out that a key implication of nanotechnology is that it will bring the economics of software to hardware, i.e., to physical products. Software prices are deflating even more quickly than hardware.

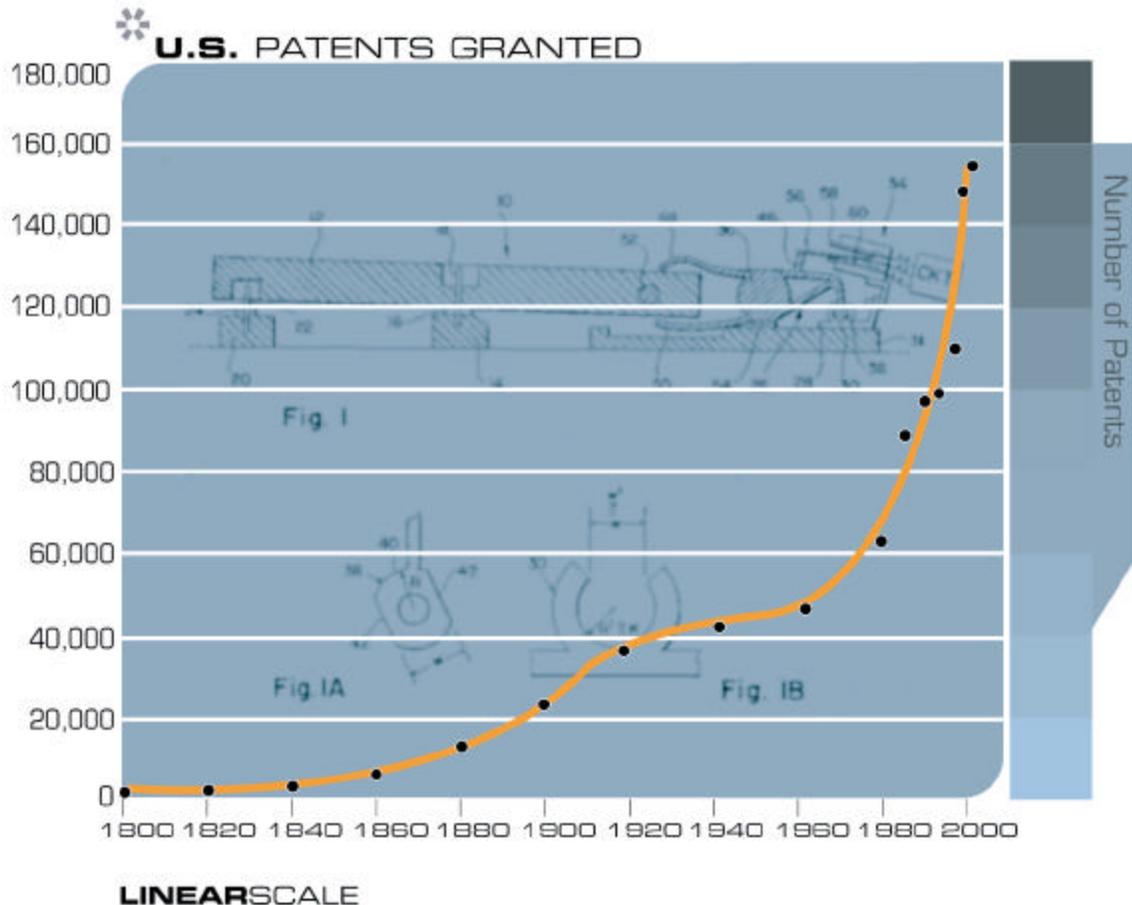
**Software Price-Performance Has Also
Improved at an Exponential Rate**
(Example: Automatic Speech Recognition Software)

	1985	1995	2000
Price	\$5,000	\$500	\$50
Vocabulary Size (# words)	1,000	10,000	100,000
Continuous Speech?	No	No	Yes
User Training Required (Minutes)	180	60	5
Accuracy	Poor	Fair	Good



Current economic policy is based on outdated models that include energy prices, commodity prices, and capital investment in plant and equipment as key driving factors, but do not adequately model the size of technology, bandwidth, MIPS, megabytes, intellectual property, knowledge, and other increasingly vital (and increasingly increasing) constituents that are driving the economy.

Another indication of the law of accelerating returns in the exponential growth of human knowledge, including intellectual property. If we look at the development of intellectual property within the nanotechnology field, we see even more rapid growth.

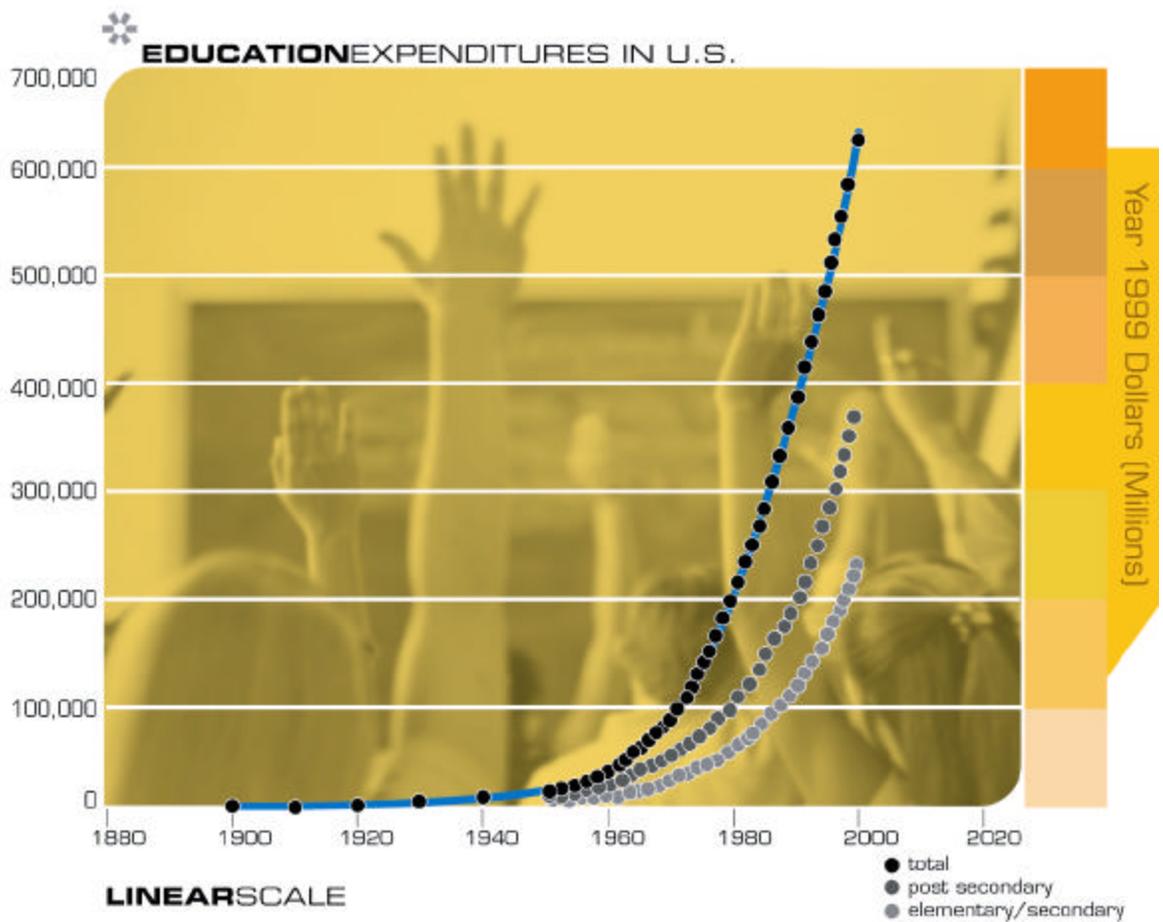


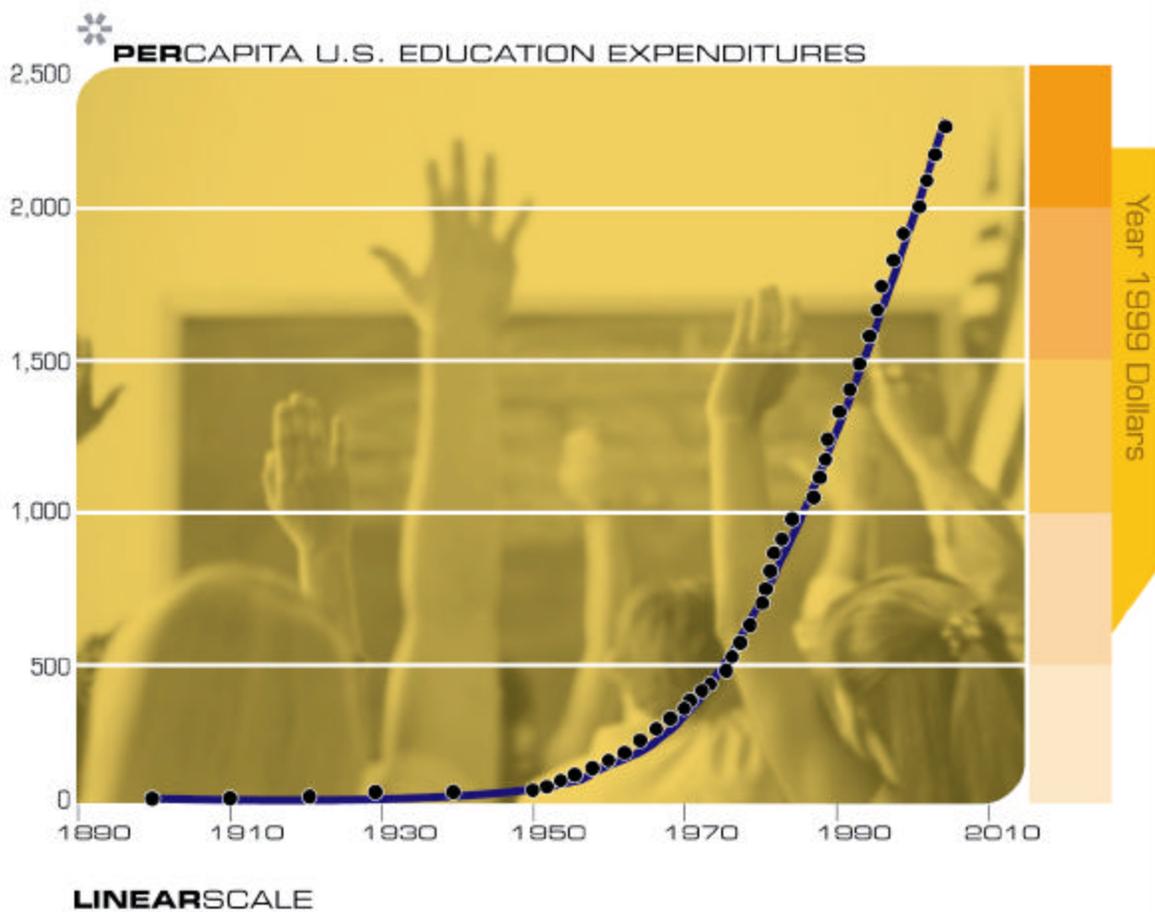
None of this means that cycles of recession will disappear immediately. Indeed there is a current economic slowdown and a technology-sector recession. The economy still has some of the underlying dynamics that historically have caused cycles of recession, specifically excessive commitments such as over-investment, excessive capital intensive projects and the overstocking of inventories. However, the rapid dissemination of information, sophisticated forms of online procurement, and increasingly transparent markets in all industries have diminished the impact of this cycle. So “recessions” are likely to have less direct impact on our standard of living. The underlying long-term growth rate will continue at a double exponential rate.

Moreover, innovation and the rate of paradigm shift are not noticeably affected by the minor deviations caused by economic cycles. All of the technologies exhibiting exponential growth shown in the above charts are continuing without losing a beat through this economic slowdown.

The overall growth of the economy reflects completely new forms and layers of wealth and value that did not previously exist, or least that did not previously constitute a significant portion of the economy (but do now): new forms of nanoparticle-based materials, genetic information, intellectual property, communication portals, web sites, bandwidth, software, data bases, and many other new technology-based categories.

Another implication of the law of accelerating returns is exponential growth in education and learning. Over the past 120 years, we have increased our investment in K-12 education (per student and in constant dollars) by a factor of ten. We have a one hundred fold increase in the number of college students. Automation started by amplifying the power of our muscles, and in recent times has been amplifying the power of our minds. Thus, for the past two centuries, automation has been eliminating jobs at the bottom of the skill ladder while creating new (and better paying) jobs at the top of the skill ladder. So the ladder has been moving up, and thus we have been exponentially increasing investments in education at all levels.





The Deeply Intertwined Promise and Peril of Nanotechnology and Related Advanced Technologies

Technology has always been a double-edged sword, bringing us longer and healthier life spans, freedom from physical and mental drudgery, and many new creative possibilities on the one hand, while introducing new and salient dangers on the other. Technology empowers both our creative and destructive natures. Stalin's tanks and Hitler's trains used technology. We still live today with sufficient nuclear weapons (not all of which appear to be well accounted for) to end all mammalian life on the planet. Bioengineering is in the early stages of enormous strides in reversing disease and aging processes. However, the means and knowledge will soon exist in a routine college bioengineering lab (and already exists in more sophisticated labs) to create unfriendly

pathogens more dangerous than nuclear weapons. As technology accelerates towards the full realization of biotechnology, nanotechnology and “strong” AI (artificial intelligence at human levels and beyond), we will see the same intertwined potentials: a feast of creativity resulting from human intelligence expanded many-fold combined with many grave new dangers.

Consider unrestrained nanobot replication. Nanobot technology requires billions or trillions of such intelligent devices to be useful. The most cost-effective way to scale up to such levels is through self-replication, essentially the same approach used in the biological world. And in the same way that biological self-replication gone awry (i.e., cancer) results in biological destruction, a defect in the mechanism curtailing nanobot self-replication would endanger all physical entities, biological or otherwise. I address below steps we can take to address this grave risk, but we cannot have complete assurance in any strategy that we devise today.

Other primary concerns include “who is controlling the nanobots?” and “who are the nanobots talking to?” Organizations (e.g., governments, extremist groups) or just a clever individual could put trillions of undetectable nanobots in the water or food supply of an individual or of an entire population. These “spy” nanobots could then monitor, influence, and even control our thoughts and actions. In addition to introducing physical spy nanobots, existing nanobots could be influenced through software viruses and other software “hacking” techniques. When there is software running in our brains, issues of privacy and security will take on a new urgency.

My own expectation is that the creative and constructive applications of this technology will dominate, as I believe they do today. However, I believe we need to invest more heavily in developing specific defensive technologies. As I address further below, we are at this stage today for biotechnology, and will reach the stage where we need to directly implement defensive technologies for nanotechnology during the late teen years of this century.

If we imagine describing the dangers that exist today to people who lived a couple of hundred years ago, they would think it mad to take such risks. On the other hand, how many people in the year 2000 would really want to go back to the short, brutish, disease-filled, poverty-stricken, disaster-prone lives that 99 percent of the human race struggled through a couple of centuries ago? We may romanticize the past, but up until fairly recently, most of humanity lived extremely fragile lives where one all-too-common misfortune could spell disaster. Substantial portions of our species still live in this precarious way, which is at least one reason to continue technological progress and the economic enhancement that accompanies it.

People often go through three stages in examining the impact of future technology: awe and wonderment at its potential to overcome age old problems; then a sense of dread at a new set of grave dangers that accompany these new technologies; followed, finally and hopefully, by the realization that the only viable and responsible path is to set a careful course that can realize the promise while managing the peril.

This congressional hearing was partly inspired by Bill Joy's cover story for *Wired* magazine, *Why The Future Doesn't Need Us*. Bill Joy, cofounder of Sun Microsystems and principal developer of the Java programming language, has recently taken up a personal mission to warn us of the impending dangers from the emergence of self-replicating technologies in the fields of genetics, nanotechnology, and robotics, which he aggregates under the label "GNR." Although his warnings are not entirely new, they have attracted considerable attention because of Joy's credibility as one of our leading technologists. It is reminiscent of the attention that George Soros, the currency arbitrageur and arch capitalist, received when he made vaguely critical comments about the excesses of unrestrained capitalism .

Joy's concerns include genetically altered designer pathogens, followed by self-replicating entities created through nanotechnology. And if we manage to survive these first two perils, we will encounter robots whose intelligence will rival and ultimately exceed our own. Such robots may make great assistants, but who's to say that we can count on them to remain reliably friendly to mere humans?

Although I am often cast as the technology optimist who counters Joy's pessimism, I do share his concerns regarding self-replicating technologies; indeed, I played a role in bringing these dangers to Bill's attention. In many of the dialogues and forums in which I have participated on this subject, I end up defending Joy's position with regard to the feasibility of these technologies and scenarios when they come under attack by commentators who I believe are being quite shortsighted in their skepticism. Even so, I do find fault with Joy's prescription: halting the advance of technology and the pursuit of knowledge in broad fields such as nanotechnology.

In his essay, Bill Joy eloquently described the plagues of centuries past and how new self-replicating technologies, such as mutant bioengineered pathogens and "nanobots" run amok, may bring back long-forgotten pestilence. Indeed these are real dangers. It is also the case, which Joy acknowledges, that it has been technological advances, such as antibiotics and improved sanitation, which have freed us from the prevalence of such plagues. Suffering in the world continues and demands our steadfast attention. Should we tell the millions of people afflicted with cancer and other devastating conditions that we are canceling the development of all bioengineered treatments because there is a risk that these same technologies may someday be used for malevolent purposes? Having asked the rhetorical question, I realize that there is a movement to do exactly that, but I think most people would agree that such broad-based relinquishment is not the answer.

The continued opportunity to alleviate human distress is one important motivation for continuing technological advancement. Also compelling are the already apparent economic gains I discussed above that will continue to hasten in the decades ahead. The continued acceleration of many intertwined technologies are roads paved with gold (I use the plural here because technology is clearly not a single path). In a competitive environment, it is an economic imperative to go down these roads. Relinquishing

technological advancement would be economic suicide for individuals, companies, and nations.

The Relinquishment Issue

This brings us to the issue of relinquishment, which is Bill Joy's most controversial recommendation and personal commitment. I do feel that relinquishment at the right level is part of a responsible and constructive response to these genuine perils. The issue, however, is exactly this: at what level are we to relinquish technology?

Ted Kaczynski would have us renounce all of it. This, in my view, is neither desirable nor feasible, and the futility of such a position is only underscored by the senselessness of Kaczynski's deplorable tactics. There are other voices, less reckless than Kaczynski, who are nonetheless arguing for broad-based relinquishment of technology. Bill McKibben, the environmentalist who was one of the first to warn against global warming, takes the position that "environmentalists must now grapple squarely with the idea of a world that has enough wealth and enough technological capability, and should not pursue more." In my view, this position ignores the extensive suffering that remains in the human world, which we will be in a position to alleviate through continued technological progress.

Another level would be to forego certain fields -- nanotechnology, for example -- that might be regarded as too dangerous. But such sweeping strokes of relinquishment are equally untenable. As I pointed out above, nanotechnology is simply the inevitable end result of the persistent trend towards miniaturization that pervades all of technology. It is far from a single centralized effort, but is being pursued by a myriad of projects with many diverse goals.

One observer wrote:

"A further reason why industrial society cannot be reformed. . . is that modern technology is a unified system in which all parts are dependent on one another. You can't get rid of the "bad" parts of technology and retain only the "good" parts. Take modern medicine, for example. Progress in medical science depends on progress in chemistry, physics, biology, computer science and other fields. Advanced medical treatments require expensive, high-tech equipment that can be made available only by a technologically progressive, economically rich society. Clearly you can't have much progress in medicine without the whole technological system and everything that goes with it."

The observer I am quoting is, again, Ted Kaczynski. Although one will properly resist Kaczynski as an authority, I believe he is correct on the deeply entangled nature of the benefits and risks. However, Kaczynski and I clearly part company on our overall assessment on the relative balance between the two. Bill Joy and I have dialogued on this issue both publicly and privately, and we both believe that technology will and should

progress, and that we need to be actively concerned with the dark side. If Bill and I disagree, it's on the granularity of relinquishment that is both feasible and desirable.

Abandonment of broad areas of technology will only push them underground where development would continue unimpeded by ethics and regulation. In such a situation, it would be the less-stable, less-responsible practitioners (e.g., terrorists) who would have all the expertise.

I do think that relinquishment at the right level needs to be part of our ethical response to the dangers of 21st century technologies. One constructive example of this is the proposed ethical guideline by the Foresight Institute, founded by nanotechnology pioneer Eric Drexler, that nanotechnologists agree to relinquish the development of physical entities that can self-replicate in a natural environment. Another is a ban on self-replicating physical entities that contain their own codes for self-replication. In what nanotechnologist Ralph Merkle calls the "broadcast architecture," such entities would have to obtain such codes from a centralized secure server, which would guard against undesirable replication. I discuss these guidelines further below.

The broadcast architecture is impossible in the biological world, which represents at least one way in which nanotechnology can be made safer than biotechnology. In other ways, nanotech is potentially more dangerous because nanobots can be physically stronger than protein-based entities and more intelligent. It will eventually be possible to combine the two by having nanotechnology provide the codes within biological entities (replacing DNA), in which case biological entities can use the much safer broadcast architecture. I comment further on the strengths and weaknesses of the broadcast architecture below.

As responsible technologies, our ethics should include such "fine-grained" relinquishment, among other professional ethical guidelines. Other protections will need to include oversight by regulatory bodies, the development of technology-specific "immune" responses, as well as computer assisted surveillance by law enforcement organizations. Many people are not aware that our intelligence agencies already use advanced technologies such as automated word spotting to monitor a substantial flow of telephone conversations. As we go forward, balancing our cherished rights of privacy with our need to be protected from the malicious use of powerful 21st century technologies will be one of many profound challenges. This is one reason that such issues as an encryption "trap door" (in which law enforcement authorities would have access to otherwise secure information) and the FBI "Carnivore" email-snooping system have been controversial, although these controversies have abated since 9-11-2001.

As a test case, we can take a small measure of comfort from how we have dealt with one recent technological challenge. There exists today a new form of fully nonbiological self replicating entity that didn't exist just a few decades ago: the computer virus. When this form of destructive intruder first appeared, strong concerns were voiced that as they became more sophisticated, software pathogens had the potential to destroy the computer network medium they live in. Yet the "immune system" that has evolved in

response to this challenge has been largely effective. Although destructive self-replicating software entities do cause damage from time to time, the injury is but a small fraction of the benefit we receive from the computers and communication links that harbor them. No one would suggest we do away with computers, local area networks, and the Internet because of software viruses.

One might counter that computer viruses do not have the lethal potential of biological viruses or of destructive nanotechnology. This is not always the case; we rely on software to monitor patients in critical care units, to fly and land airplanes, to guide intelligent weapons in our current campaign in Iraq, and other “mission-critical” tasks. To the extent that this is true, however, this observation only strengthens my argument. The fact that computer viruses are not usually deadly to humans only means that more people are willing to create and release them. It also means that our response to the danger is that much less intense. Conversely, when it comes to self-replicating entities that are potentially lethal on a large scale, our response on all levels will be vastly more serious, as we have seen since 9-11.

I would describe our response to software pathogens as effective and successful. Although they remain (and always will remain) a concern, the danger remains at a nuisance level. Keep in mind that this success is in an industry in which there is no regulation, and no certification for practitioners. This largely unregulated industry is also enormously productive. One could argue that it has contributed more to our technological and economic progress than any other enterprise in human history. I discuss the issue of regulation further below.

Development of Defensive Technologies and the Impact of Regulation

Joy’s treatise is effective because he paints a picture of future dangers as if they were released on today’s unprepared world. The reality is that the sophistication and power of our defensive technologies and knowledge will grow along with the dangers. When we have “gray goo” (unrestrained nanobot replication), we will also have “blue goo” (“police” nanobots that combat the “bad” nanobots). The story of the 21st century has not yet been written, so we cannot say with assurance that we will successfully avoid all misuse. But the surest way to prevent the development of the defensive technologies would be to relinquish the pursuit of knowledge in broad areas. We have been able to largely control harmful software virus replication because the requisite knowledge is widely available to responsible practitioners. Attempts to restrict this knowledge would have created a far less stable situation. Responses to new challenges would have been far slower, and it is likely that the balance would have shifted towards the more destructive applications (e.g., software viruses).

The challenge most immediately in front of us is not self-replicating nanotechnology, but rather self-replicating biotechnology. The next two decades will be the golden age of biotechnology, whereas the comparable era for nanotechnology will follow in the 2020s and beyond. We are now in the early stages of a transforming technology based on the intersection of biology and information science. We are learning

the “software” methods of life and disease processes. By reprogramming the information processes that lead to and encourage disease and aging, we will have the ability to overcome these afflictions. However, the same knowledge can also empower a terrorist to create a bioengineered pathogen.

As we compare the success we have had in controlling engineered software viruses to the coming challenge of controlling engineered biological viruses, we are struck with one salient difference. As I noted above, the software industry is almost completely unregulated. The same is obviously not the case for biotechnology. A bioterrorist does not need to put his “innovations” through the FDA. However, we do require the scientists developing the defensive technologies to follow the existing regulations, which slow down the innovation process at every step. Moreover, it is impossible, under existing regulations and ethical standards, to test defenses to bioterrorist agents. There is already extensive discussion to modify these regulations to allow for animal models and simulations to replace infeasible human trials. This will be necessary, but I believe we will need to go beyond these steps to accelerate the development of vitally needed defensive technologies.

For reasons I have articulated above, stopping these technologies is not feasible, and pursuit of such broad forms of relinquishment will only distract us from the vital task in front of us. In terms of public policy, the task at hand is to rapidly develop the defensive steps needed, which include ethical standards, legal standards, and defensive technologies. It is quite clearly a race. As I noted, in the software field, the defensive technologies have remained a step ahead of the offensive ones. With the extensive regulation in the medical field slowing down innovation at each stage, we cannot have the same confidence with regard to the abuse of biotechnology.

In the current environment, when one person dies in gene therapy trials, there are congressional investigations and all gene therapy research comes to a temporary halt. There is a legitimate need to make biomedical research as safe as possible, but our balancing of risks is completely off. The millions of people who desperately need the advances to be made available by gene therapy and other breakthrough biotechnology advances appear to carry little political weight against a handful of well-publicized casualties from the inevitable risks of progress.

This equation will become even more stark when we consider the emerging dangers of bioengineered pathogens. What is needed is a change in public attitude in terms of tolerance for needed risk.

Hastening defensive technologies is absolutely vital to our security. We need to streamline regulatory procedures to achieve this. However, we also need to greatly increase our investment explicitly in the defensive technologies. In the biotechnology field, this means the rapid development of antiviral medications. We will not have time to develop specific countermeasures for each new challenge that comes along. We are close to developing more generalized antiviral technologies, and these need to be accelerated.

I have addressed here the issue of biotechnology because that is the threshold and challenge that we now face. The comparable situation will exist for nanotechnology once replication of nano-engineered entities has been achieved. As that threshold comes closer, we will then need to invest specifically in the development of defensive technologies, including the creation of a nanotechnology-based immune system. Bill Joy and other observers have pointed out that such an immune system would itself be a danger because of the potential of “autoimmune” reactions (i.e., the immune system using its powers to attack the world it is supposed to be defending).

However, this observation is not a compelling reason to avoid the creation of an immune system. No one would argue that humans would be better off without an immune system because of the possibility of auto immune diseases. Although the immune system can itself be a danger, humans would not last more than a few weeks (barring extraordinary efforts at isolation) without one. The development of a technological immune system for nanotechnology will happen even without explicit efforts to create one. We have effectively done this with regard to software viruses. We created a software virus immune system not through a formal grand design project, but rather through our incremental responses to each new challenge. We can expect the same thing will happen as challenges from nanotechnology based dangers emerge. The point for public policy will be to specifically invest in these defensive technologies.

It is premature today to develop specific defensive nanotechnologies since we can only have a general idea of what we are trying to defend against. It would be similar to the engineering world creating defenses against software viruses before the first one had been created. However, there is already fruitful dialogue and discussion on anticipating this issue, and significantly expanded investment in these efforts is to be encouraged.

As I mentioned above, the Foresight Institute, for example, has devised a set of ethical standards and strategies for assuring the development of safe nanotechnology. These guidelines include:

- “Artificial replicators must not be capable of replication in a natural, uncontrolled environment.”
- “Evolution within the context of a self-replicating manufacturing system is discouraged.”
- “MNT (molecular nanotechnology) designs should specifically limit proliferation and provide traceability of any replicating systems.”
- “Distribution of molecular manufacturing development capability should be restricted whenever possible, to responsible actors that have agreed to the guidelines. No such restriction need apply to end products of the development process.”

Other strategies that the Foresight Institute has proposed include:

- Replication should require materials not found in the natural environment.

- Manufacturing (replication) should be separated from the functionality of end products. Manufacturing devices can create end products, but cannot replicate themselves, and end products should have no replication capabilities.
- Replication should require replication codes that are encrypted, and time limited. The broadcast architecture mentioned earlier is an example of this recommendation.

These guidelines and strategies are likely to be effective with regarding to preventing accidental release of dangerous self-replicating nanotechnology entities. The situation with regard to intentional design and release of such entities is more complex and more challenging. We can anticipate approaches that would have the potential to defeat each of these layers of protections by a sufficiently determined and destructive opponent.

Take, for example, the broadcast architecture. When properly designed, each entity is unable to replicate without first obtaining replication codes. These codes are not passed on from one replication generation to the next. However, a modification to such a design could bypass the destruction of the replication codes and thereby pass them on to the next generation. To overcome that possibility, it has been recommended that the memory for the replication codes be limited to only a subset of the full replication code so that there is insufficient memory to pass the codes along. However, this guideline could be defeated by expanding the size of the replication code memory to incorporate the entire code. Another protection that has been suggested is to encrypt the codes and to build in protections such as time expiration limitations in the decryption systems. However, we can see the ease with which protections against unauthorized replications of intellectual property such as music files has been defeated. Once replication codes and protective layers are stripped away, the information can be replicated without these restrictions.

My point is not that protection is impossible. Rather, we need to realize that any level of protection will only work to a certain level of sophistication. The “meta” lesson here is that we will need to continue to advance the defensive technologies, and keep them one or more steps ahead of the destructive technologies. We have seen analogies to this in many areas, including technologies for national defense, as well as our largely successful efforts to combat software viruses, that I alluded to above.

What we can do today with regard to the critical challenge of self-replication in nanotechnology is to continue the type of effective study that the Foresight Institute has initiated. With the human genome project, three to five percent of the budgets were devoted to the ethical, legal, and social implications (ELSI) of the technology. A similar commitment for nanotechnology would be appropriate and constructive.

Technology will remain a double-edged sword, and the story of the 21st century has not yet been written. It represents vast power to be used for all humankind’s purposes. We have no choice but to work hard to apply these quickening technologies to

advance our human values, despite what often appears to be a lack of consensus on what those values should be.